

Testing Light Dark Matter Coannihilation With Fixed-Target Experiments

Eder Izaguirre,^{1,*} Yonatan Kahn,^{2,†} Gordan Krnjaic,^{3,‡} and Matthew Moschella^{2,§}

¹Brookhaven National Laboratory, Upton, NY

²Princeton University, Princeton, NJ

³Fermi National Accelerator Laboratory, Batavia, IL

In this paper, we introduce a novel program of fixed-target searches for thermal-origin Dark Matter (DM), which couples inelastically to the Standard Model. Since the DM only interacts by transitioning to a heavier state, freeze-out proceeds via coannihilation and the unstable heavier state is depleted at later times. For sufficiently large mass splittings, direct detection is kinematically forbidden and indirect detection is impossible, so this scenario can only be tested with accelerators. Here we propose new searches at proton and electron beam fixed-target experiments to probe sub-GeV coannihilation, exploiting the distinctive signals of up- and down-scattering as well as decay of the excited state inside the detector volume. We focus on a representative model in which DM is a pseudo-Dirac fermion coupled to a hidden gauge field (dark photon), which kinetically mixes with the visible photon. We define theoretical targets in this framework and determine the existing bounds by reanalyzing results from previous experiments. We find that LSND, E137, and BaBar data already place strong constraints on the parameter space consistent with a thermal freeze-out origin, and that future searches at Belle II and MiniBooNE, as well as recently-proposed fixed-target experiments such as LDMX and BDX, can cover nearly all remaining gaps. We also briefly comment on the discovery potential for proposed beam dump and neutrino experiments which operate at much higher beam energies.

I. INTRODUCTION

Understanding the particle nature of Dark Matter (DM) is among the highest priorities in all of physics. Perhaps the most popular DM candidate to date has been the Weakly Interacting Massive Particle (WIMP), which is charged under the electroweak force and naturally yields the observed cosmological abundance via thermal freeze-out (see [1] for a review). However, decades of null results from direct detection, indirect detection, and collider searches have cast doubt on this paradigm and motivated many alternative possibilities [2, 3].

Nonetheless, the thermal freeze-out mechanism remains compelling even if DM is not a WIMP. First and foremost, thermal DM is largely UV insensitive; its abundance is determined by the DM particle properties and is unaffected by the details of earlier, unknown cosmological epochs (e.g. reheating). Furthermore, unlike nonthermal production mechanisms, which can accommodate DM masses anywhere between 10^{-22} eV – $10 M_\odot$ (!), thermal DM is only viable between ~ 5 keV – 100 TeV, and is therefore more predictive. Dark matter masses outside this window are either too hot for acceptable structure formation [4] or violate perturbative unitarity [5]. Finally, achieving the observed abundance requires a minimum interaction rate between DM and the Standard Model (SM), which provides a clear target for discovery or falsification. Thus, there is ample motivation to identify and study every viable realization of this mechanism.

One simple way to completely eliminate the tension between a thermal origin and experimental limits, in particular

those from direct detection experiments, is for the DM to couple *inelastically* to SM particles [6]. In this class of models, the halo DM species χ_1 is the lightest stable particle in the dark sector and interacts with the SM only by transitioning to a slightly heavier state χ_2 . This class of models has several appealing features:

- **Large Viable Couplings:** If the inelastic interaction with the SM also determines the leading annihilation process, the relic abundance arises dominantly through $\chi_1 \chi_2 \rightarrow \text{SM}$, a process dubbed *coannihilation* [7]. Since the heavier χ_2 population is Boltzmann-suppressed during freeze-out, the requisite annihilation rate must compensate for this penalty.¹ Thus, the coannihilation cross section always satisfies $\sigma v \geq 3 \times 10^{-26} \text{cm}^3 \text{s}^{-1}$.
- **Indirect Detection Shuts Off:** Since the heavier state is unstable, its population is fully depleted at low temperatures, so there are no remaining coannihilation partners for the χ_1 . This effect turns off possible late-time indirect detection signals and alleviates the bound from cosmic microwave background (CMB) power injection, which otherwise naively rules out thermal DM for masses below ~ 10 GeV for *s*-wave annihilation [9].²
- **Direct Detection Forbidden:** For a nonrelativistic halo particle scattering off a stationary SM target, the energy available to upscatter into the heavier state is $\sim \mu v^2$, where μ is the reduced mass of the DM-target system.

* eder@bnl.gov

† ykahn@princeton.edu

‡ krnjaic@fnal.gov

§ moschella@princeton.edu

¹ For a general scenario where coannihilation proceeds without inelastic couplings, but an analogous enhanced thermal cross section appears, see Ref. [8].

² Another way to evade this bound is “forbidden DM” [10].

Thus, for typical halo velocities $v \sim 10^{-3}$, direct detection off of nuclei is forbidden if the mass splitting exceeds $\gtrsim 100$ keV [6]. There is the possibility of loop-level induced elastic scattering off of electrons, which may be very relevant for new proposals for electron direct detection [11–16]. However, we leave this question for a future study.

Since direct and indirect detection search strategies are not available, testing thermal coannihilation fundamentally requires accelerator-based techniques. For DM masses near the weak scale (few GeV – 100s of GeV), Refs. [17–19] proposed LHC and B-factory searches with sensitivity to thermal coannihilation over a wide range of masses and splittings. However, for DM masses below the GeV scale, these searches become ineffective and new tools are required to fully test the keV–GeV mass range over which thermal coannihilation remains viable, yet unexplored.

In this paper, we fill this large gap by recasting and proposing a series of fixed-target searches for both electron and proton beam facilities. In both cases, an incident beam impinges on a fixed target and produces a boosted $\chi_1\chi_2$ pair. Depending on the experimental setup, this system can give rise to a variety of possible signals:

- **Beam Dumps:** The $\chi_1\chi_2$ pair can be produced radiatively through the “dark bremsstrahlung” reactions $pN \rightarrow pN\chi_1\chi_2$ at proton beam dumps, or $eN \rightarrow eN\chi_1\chi_2$ at electron beam dumps, where N is a nuclear target. At proton beam dumps, the DM system can also be produced from meson decays $\pi^0 \rightarrow \gamma\chi_1\chi_2$ and $\eta \rightarrow \gamma\chi_1\chi_2$, which can be advantageous for low DM masses. Once produced, the $\chi_{1,2}$ can reach a downstream detector and scatter off electrons or nucleons to induce an observable recoil signature. Alternatively, the unstable χ_2 can also decay in flight as it passes through the detector via $\chi_2 \rightarrow \chi_1 e^+ e^-$, thereby depositing an observable signal. These processes are depicted schematically in Fig. 2.
- **Electron Missing Energy/Momentum:** As above, the $\chi_1\chi_2$ pair is produced in $e^-N \rightarrow e^-N\chi_1\chi_2$ “dark bremsstrahlung”, but the production takes place in a thin target embedded in a detector. The principal observable signature in this context is the recoiling final-state electron. If this electron emerges having lost most of its incident beam energy, with no additional activity deposited in a downstream detector, this process can be sensitive to DM production. As above, a χ_2 decaying inside the detector provides an additional potential signature. This process is depicted schematically in Fig. 3.

We will show that existing data already rules out large portions of the direct coannihilation parameter space. Moreover, the dedicated searches we propose which exploit the unique signals from this class of models can significantly improve on the current levels of sensitivity. Crucially, as shown in Fig. 4, the χ_2 has a macroscopic decay length over nearly all of the parameter space we are interested in, and detecting the electromagnetic energy deposited by a χ_2 decay in the detector




FIG. 1. Leading order diagram for $\chi_1\chi_2 \rightarrow f^+f^-$ coannihilation, which sets the DM relic abundance in the $m_{A'} > m_{1,2}$ regime.

— a signal not present in elastic DM models — is sufficient to cover large portions of the thermal relic curve.³ Indeed, we show in Sec. III C that the sensitivity to the decay of the excited state χ_2 typically dominates over scattering channels at experiments with beam energies below 10 GeV. Thus, the prospects are excellent for dedicated experiments sensitive to these signatures, and we show that current and planned experiments can confirm or rule out nearly the entire mass range allowed for thermal coannihilating DM. We note a related study from Ref. [21] that investigated previous limits from fixed target facilities in the context of a supersymmetric hidden sector. Recently, Ref. [22] proposed the signal of $\chi_1 \rightarrow \chi_2$ upscattering followed by χ_2 decay for the case of (boosted) astrophysical DM, as opposed to DM produced in a beam dump.

This paper is organized as follows. In Sec. II we introduce a representative renormalizable model featuring a pseudo-Dirac fermion DM coupled to a kinetically-mixed dark photon, where the abundance of the former arises from thermal coannihilation. In Sec. III we make some general comments on the production modes and detection signals at proton- and electron-beam fixed-target experiments. In Sec. IV we describe our methods for reinterpreting existing data from LSND [25] and E137 [26], and we discuss projections for the future data at MiniBooNE [27] and the proposed BDX [28] and LDMX [23] experiments. In Sec. V we discuss the bounds and reach projections from these experiments in the context of the thermal coannihilating inelastic DM. Finally, in Sec. VI we offer some concluding remarks. Additional details on the matrix elements and Monte Carlo methods used in determining the thermal parameter space and in our simulations can be found in the Appendices.

II. SUB-GEV THERMAL COANNIHILATION

In this section, we describe a class of models of coannihilating DM: DM that couples inelastically to the SM through a kinetically-mixed dark photon. We detail the early universe cosmology and freeze out of the model, as well as introduce a useful parametrization of the parameters of the model in

³ A similar search strategy was proposed for the decay of long-lived scalars in Ref. [20].




FIG. 2. Inelastic DM production at electron and proton beam dump experiments via dark bremsstrahlung and meson decay. The resulting χ_1, χ_2 pair can give rise to a number of possible signatures in the detector: χ_2 can decay inside the fiducial volume to deposit electromagnetic energy; both χ_1 and χ_2 can scatter off detector targets T and impart visible recoil energies to these particles; or χ_1 can upscatter into χ_2 , which can then decay promptly inside the detector to deposit a visible signal.




FIG. 3. Inelastic DM production at electron beam fixed-target missing energy/momentum experiments. **Left:** Setup for an LDMX style missing momentum experiment [2, 23] in which a (\sim few GeV) beam electron produces DM in a thin target (\ll radiation length) and thereby loses a large fraction of its incident energy. The emerging lower energy electron passes through tracker material and registers as a signal event if there is no additional energy deposited in the ECAL/HCAL system downstream, which serves primarily to veto SM activity. **Right:** Setup for an NA64 style experiment in which the beam (typically at higher energies, \sim 30 GeV) produces the DM system by interacting with an instrumented, active target volume [24]. As with LDMX, the instrumented region serves to verify that the beam electron has abruptly lost most of its energy and that there is no additional SM activity downstream.

which the thermal target is largely an invariant under variation of couplings and of mass hierarchies.

A. Mediator Model Building

Unlike weak-scale WIMPs, which realize successful freeze-out with only SM gauge interactions, sub-GeV DM is overproduced in the absence of light ($\ll m_Z$) new mediators to generate a sufficiently large annihilation rate [29, 30]. To avoid detection thus far, such mediators must be neutral under the SM and couple non-negligibly to visible particles.

If SM particles are neutral under the new interaction, a

renormalizable model (without additional fields) requires the mediator to interact with the SM through the hypercharge, Higgs, or lepton portals

$$B_{\mu\nu} \ , \ H^\dagger H \ , \ LH, \quad (1)$$

for vector, scalar, and fermionic mediators, respectively. However, coupling a fermionic mediator to the lepton portal requires additional model building⁴ and scalar mediators, which mix with the Higgs are ruled out for predictive models in which DM annihilates directly to SM final states (see

⁴ A fermionic mediator coupled to the lepton portal requires additional




FIG. 4. Rest frame χ_2 decay length for various mass splittings $\Delta \equiv m_2 - m_1$, as a function of the χ_2 mass m_2 . The vertical axis is normalized to values of the dimensionless coupling $y \equiv \epsilon^2 \alpha_D (\frac{m_1}{m_{A'}})^4$ which achieve the observed relic density; see Sec. II for more details.

Sec. IIC and [31] for a discussion of this issue), so we restrict our attention to abelian vector mediators; a nonabelian field strength is not gauge invariant, so kinetic mixing is forbidden.

Alternatively, the mediator could couple directly to SM particles if both dark and visible matter are charged under the same gauge group. In the absence of additional fields, anomaly cancellation restricts the possible choices to be

$$U(1)_{B-L}, U(1)_{\ell_i - \ell_j}, U(1)_{3B - \ell_i}, \quad (2)$$

and linear combinations thereof. In most contexts, the relevant phenomenology in fixed-target searches is qualitatively similar to the vector portal scenario, so below we will ignore these possibilities without loss of essential generality. We note, however, that viable models for both protophobic [32] and protophilic [33] mediators exist, so the complementarity provided by both proton- and electron-beam experiments is highly advantageous.

B. Representative Model

Our representative dark sector contains a 4-component fermion ψ that transforms under a hidden abelian $U(1)_D$ gauge group. The fermion couples to a vector mediator A' as

$$\mathcal{L} = i\bar{\psi}\not{D}\psi + M\bar{\psi}\psi + \lambda\phi\bar{\psi}^c\psi + h.c., \quad (3)$$

model building to simultaneously achieve a thermal contact through this interaction and yield viable neutrino textures; the coupling to the mediator must be suppressed by neutrino masses, so it is generically difficult for the interaction rate to exceed Hubble expansion.

where ϕ is a $U(1)_D$ symmetry breaking scalar whose vacuum expectation value v_D gives A' a nonzero mass $m_{A'} \sim g_D v_D$ and gives ψ a Majorana mass $\sim \lambda v_D$. Diagonalizing the resulting Dirac and Majorana masses gives rise to fermion mass eigenstates $\chi_{1,2}$ with a small mass splitting $\Delta \equiv m_2 - m_1$ and an off-diagonal coupling to A' ,

$$\mathcal{L} \supset g_D A'_\mu \bar{\chi}_2 \gamma^\mu \chi_1 + h.c., \quad (4)$$

where $g_D \equiv \sqrt{4\pi\alpha_D}$ is the dark coupling constant. Note that it is technically natural to have $\Delta \ll M$ since the Majorana mass breaks the global symmetry associated with ψ number.⁵

This sector can interact with the SM through a renormalizable and gauge-invariant kinetic mixing term between $U(1)_D$ and $U(1)_Y$ gauge fields,

$$\mathcal{L} \supset \frac{\epsilon}{2 \cos \theta_W} F'_{\mu\nu} B^{\mu\nu} = \epsilon F'_{\mu\nu} F^{\mu\nu} + \epsilon \tan \theta_W F'_{\mu\nu} Z^{\mu\nu}, \quad (5)$$

where $\epsilon \ll 1$ is the kinetic mixing parameter and $F'_{\mu\nu}$ and $B_{\mu\nu}$ are respectively the dark and hypercharge field strength tensors and the kinetic mixing interaction has been written in terms of the SM mass eigenstates A and Z after electroweak symmetry breaking. Diagonalizing the kinetic terms in Eq. (5) and rescaling the field strengths to restore canonical normalization induce a coupling between A' and the SM fermions [34]. To leading order in ϵ , the A' -SM interaction becomes

$$e A_\mu J_{EM}^\mu \rightarrow e (A_\mu + \epsilon A'_\mu) J_{EM}^\mu, \quad (6)$$

where J_{EM}^μ is the SM electromagnetic current and all charged fermions acquire millicharges under $U(1)_D$. There is also an analogous A' interaction with the SM neutral current that arises from $A' - Z$ mixing, but in our mass range of interest, $m_{A'} \ll m_Z$, the mixing parameter is suppressed by an additional factor of $(m_{A'}/m_Z)^2$ [35–38], so we neglect this interaction for the remainder of paper.

C. Direct Coannihilation vs. Secluded Annihilation

In the hot early universe ($T \gg m_i, m_{A'}$), all dark species are in chemical and kinetic equilibrium with the SM radiation bath; this initial condition is guaranteed as long as the DM-SM scattering rate exceeds the Hubble expansion rate at some point during cosmic history. If $m_i > m_{A'}$, the freeze-out process is analogous to that of WIMP models. Below the freeze-out temperature $T_f \sim m_{1,2}/20$, the number densities of both species are depleted predominantly through $\chi_i \chi_i \rightarrow A' A'$ self-annihilation (which depends only on α_D), *not* coannihilation, which depends on the combination $\epsilon^2 \alpha_D$ and is greatly suppressed by comparison. Although both components undergo freeze-out separately, since χ_2 is heavier and

⁵ If, unlike the construction in Eq. (3), the Majorana masses for the two Weyl components in $\psi = (\xi, \eta^\dagger)$ are different, there is also a subleading diagonal interaction of the form $(\delta/M_D) \bar{\chi}_i A' \chi_i$, where $\delta \equiv m_\xi - m_\eta$ is the difference of Majorana masses for the interaction eigenstates. We neglect this interaction in our analysis, assuming the off-diagonal interaction dominates.

unstable, it will be depleted through downscattering and decays. Thus, up to order-one corrections, the requisite self-annihilation cross section satisfies the familiar WIMP-like requirement $\langle\sigma v\rangle \sim 3 \times 10^{-26} \text{cm}^3 \text{s}^{-1}$ in order for χ_1 to have the observed abundance at late times.

However, this *secluded* ($m_i > m_{A'}$) regime has several drawbacks. Since the self-annihilation rate for fermions is s -wave, annihilation continues to occur out of equilibrium during recombination, which ionizes newly-formed hydrogen and thereby modifies the CMB power spectrum. For a thermal annihilation rate, this bound rules out DM below ~ 10 GeV [9].⁶ Furthermore, since the secluded annihilation cross section scales as $\sigma v \sim \alpha_D^2/m_i^2$, the abundance is independent of the A' coupling to SM states, so there is no minimum interaction strength target for the DM search effort — at direct detection or accelerators, since these are sensitive to ϵ .

By contrast, in the *direct coannihilation* regime ($m_i < m_{A'}$), the two species annihilate each other via $\chi_1 \chi_2 \rightarrow f \bar{f}$, which has several compelling features:

- **Predictive:** The annihilation rate depends crucially on the mixing with the SM, $\sigma v \propto \epsilon^2 \alpha_D$, so for dark couplings that satisfy perturbative unitarity, there is a minimum value of ϵ that is compatible with thermal freeze-out. Reaching experimental sensitivity to this minimum value for each viable DM mass suffices to discover or falsify this entire class of models.
- **Large Cross Section:** Unlike secluded annihilation, which involves the annihilation of equal mass particles, direct coannihilation requires both χ_1 and χ_2 in the initial state. However, χ_2 typically becomes Boltzmann-suppressed before the nominal χ_1 freeze-out temperature $\sim m_1/20$, so the coannihilation cross section must be larger to compensate for this reduction. Thus, we generically require $\langle\sigma v\rangle \gg 3 \times 10^{-26} \text{cm}^3 \text{s}^{-1}$ to achieve the observed abundance, where the precise value grows sensitively with increasing Δ (compare the relic density curves in Fig. 6).
- **Bounded Viable Range:** Since the necessary coannihilation rate for freeze-out grows sharply as the mass splitting is increased, for sufficiently large Δ , *i.e.*, $\Delta \sim m_1$, the requisite ϵ is easily excluded independently of other model properties (e.g. by precision QED/electroweak measurements). Thus, as we will see below, this class of coannihilating models can be tested over the full viable parameter space.

Thus, for the remainder of this paper, our focus will be on the direct coannihilation scenario. For simplicity, without loss of essential generality, we will also demand that $m_{A'} > m_1 + m_2$ so that A' can decay to dark sector states — otherwise, the A' decays to SM final state, a signature for which there are

abundant ideas to cover [2]. Since $\epsilon \ll 1$, the A' branching ratio to SM states is correspondingly negligible in this regime.

Note that if we had chosen a neutral scalar mediator instead of the vector A' , it could only couple to SM fermions by mixing with the Higgs after electroweak symmetry breaking. However, such a scenario cannot viably realize thermal coannihilation below the GeV scale because the requisite Higgs-mediator mixing angle must be $\sim \mathcal{O}(1)$ to overcome the Yukawa penalty in the annihilation diagram and yield a thermal annihilation rate. Such large couplings are ruled out by Higgs coupling measurements and rare meson decay searches [31, 39]. We note that t -channel annihilation into a light scalar mediator can still realize secluded annihilation, but this process is less predictive and beyond the scope of this work.

D. Covering the Thermal Target

The goal of this paper is to compute existing bounds on the parameter space for thermal freeze-out via coannihilation and present sensitivity projections for future experiments. One major challenge of such an effort is the large dimensionality of this parameter space: each viable model point is defined by a unique choice of the inputs $\{m_1, m_{A'}, \Delta, \epsilon, \alpha_D\}$, constrained by the requirement that the coannihilation rate yields the observed DM density. Fortunately, in the nonrelativistic limit, for each choice of Δ , the coannihilation rate scales as

$$\sigma v \propto \epsilon^2 \alpha_D \frac{m_1^2}{m_{A'}^4} = \left(\epsilon^2 \alpha_D \frac{m_1^4}{m_{A'}^4} \right) \frac{1}{m_1^2} \equiv \frac{y}{m_1^2}, \quad (7)$$

which is valid for $s \simeq (m_1 + m_2)^2$ and sufficiently away from the s -channel resonance at $m_{A'} \sim m_1 + m_2$. Here we have defined the dimensionless parameter

$$y \equiv \epsilon^2 \alpha_D \left(\frac{m_1}{m_{A'}} \right)^4, \quad (8)$$

which uniquely determines the freeze-out annihilation rate for a given choice of m_1 and Δ .

The virtue of this parameterization is that the observed relic density is achieved only along a one-parameter curve $y(m_1)$ which is insensitive to the relative sizes of ϵ , α_D , or $m_1/m_{A'}$, reducing the dimensionality of the parameter space. However, the drawback is that some experiments only constrain a subset of model parameters and are not sensitive to the same combination of inputs that define the y variable in Eq. (8). For example, lepton colliders can be used to constrain ϵ as a function of $m_{A'}$ in searches for $e^+e^- \rightarrow \gamma \cancel{e}$, which can be interpreted as $e^+e^- \rightarrow \gamma A'$ production followed by prompt A' decay to invisible particles [40–42]. Since $\text{Br}(A' \rightarrow \text{invisible}) \simeq 1$, the signal event yield depends on ϵ , but is insensitive to α_D or the ratio $m_1/m_{A'}$ for larger χ_2 lifetimes.

Nonetheless, it is still possible to use this information to constrain the y variable for a conservative comparison with the thermal target. The strategy is to construct the *largest* possible y_{exp} using the experimentally determined ϵ_{exp} ,

$$y_{\text{exp}} = \epsilon_{\text{exp}}^2 \times \left[\alpha_D \left(\frac{m_1}{m_{A'}} \right)^4 \right], \quad (9)$$

⁶ If instead, the DM is a scalar and annihilates directly to SM fermions through an s -channel vector mediator, its annihilation rate is p -wave suppressed, which can evade CMB bounds.




FIG. 5. Decay reach in ϵ as a function of α_D , with other model parameters held fixed, for various fixed-target experiments. The value of ϵ corresponding to the thermal target y for each α_D is shown for comparison in black. The most conservative choice for α_D relative to the thermal target is the largest α_D consistent with perturbativity. Note that the different slope of the LSND curve at 50 MeV compared to the other experiments is a consequence of off-shell A' decay, see Sec. III A. Also note that the MiniBooNE projection from the right panel is missing from the left panel because the mass splitting on the left is 1.5 MeV, which is below the threshold to pass the experimental cut in Eq. (25). See Sec. IV for further details on computing the reach.

where the quantity in square brackets is chosen to be as large as possible while remaining consistent with both perturbative unitarity and the requirement that direct coannihilation remains the dominant annihilation process. Thus, the conservative prescription is to adopt order-one values for α_D and $m_1/m_{A'}$, while taking $m_{A'} > m_1 + m_2$ and avoiding the $m_{A'} \sim m_1 + m_2$ resonance. As an illustration, Fig. 5 shows the reach in ϵ for χ_2 decay at various experiments, demonstrating that the weakest reach with respect to the thermal target occurs at large α_D . For our numerical results in Figs. 6 and 8, we chose representative benchmarks $\alpha_D = 0.1, m_1/m_{A'} = 1/3$ where appropriate (e.g. for collider bounds). In Fig. 7 and in the discussion in Sec. V below, we show how the constraints on the parameter space change for different choices of these benchmark values, and demonstrate that our conclusions are qualitatively unchanged.

Note that for our main results, which are discussed in Sec. V and presented in Figs. 6, 7, and 8, we compute the $y(m_1)$ curve by numerically solving the full Boltzmann system for this model as described in Appendix A.

III. PRODUCTION AND DETECTION BASICS

In this section we review the basic A' production formalism in proton- and electron-beam fixed-target collisions to develop some intuition for our later numerical results. We also present the formalism for various DM detection signatures that ensue from boosted $A' \rightarrow \chi_1 \chi_2$ decays.

A. Production at Proton Beam Dumps

At proton beam dumps, $\chi_1 \chi_2$ pairs can be produced either from dark bremsstrahlung $pN \rightarrow pN A'$, $A' \rightarrow \chi_1 \chi_2$ or from the neutral meson decays $m^0 \rightarrow \gamma A'^{(*)} \rightarrow \gamma \chi_1 \chi_2$ where $m^0 = \pi^0, \eta$.⁷ The meson production rates and spectra depend on various experimental factors, but the strongest dependence is on the proton beam energy, while the detailed material properties of the beam dump can largely be ignored. At few-GeV beam energies, $N_{\pi^0} \approx N_p$ and $N_\eta \approx N_p/30$ [43] where N_p is the total flux of protons. See Ref. [44] for a detailed discussion of production modes and distributions.

If $m_{A'} < m_{m^0}$, the A' can be produced on-shell and the number of $\chi_1 \chi_2$ pairs via meson decays is given by

$$N_{\chi_1 \chi_2} = N_{m^0} \times 2\epsilon^2 \left(1 - \frac{m_{A'}^2}{m_{m^0}^2}\right)^3 \text{BR}(m^0 \rightarrow \gamma\gamma), \quad (10)$$

where N_{m^0} is the number of π^0 or η produced.⁸ In either case, there is a kinematic threshold for production through mesons: $m_1 + m_2 = 2m_1 + \Delta < m_{\pi^0} \approx 135$ MeV for pions, and $2m_1 + \Delta < m_\eta \approx 548$ MeV for etas.

⁷ At sufficiently high energies, direct A' production from quarks and gluons in the target becomes relevant, but this process is negligible for the beam energies we consider in this work.

⁸ If $m_{A'} > m_{m^0}$, the meson decay proceeds through an off-shell A' and the rate scales as $\epsilon^2 \alpha_D$ [45], in contrast to on-shell production which is independent of α_D . At low-energy experiments such as LSND with only a single production channel, off-shell production may be important, but for higher-energy experiments such as MiniBooNE with multiple production channels, on-shell production is dominant.

The rate of dark bremsstrahlung production is proportional to ϵ^2 , but varies non-trivially with the dark photon mass $m_{A'}$. Using the simulation of Ref. [44], we find that the number of A' produced varies from $\sim 10^{-4}\epsilon^2 N_p$ as $m_{A'} \rightarrow 0$ to $\sim 10^{-9}\epsilon^2 N_p$ at $m_{A'} = 3$ GeV. A' production can be enhanced through vector meson mixing near $m_{A'} = m_\rho, m_\omega$, but since this enhancement is only in a limited mass range and depends on the precise choice of $m_1/m_{A'}$, we do not attempt to model it precisely. The dark bremsstrahlung production mechanism has no mass threshold and so heavy DM can still be produced up to the beam energy, $2m_1 + \Delta < E_{\text{beam}}$.

B. Production at Electron Beams

At electron-beam experiments, mesons are no longer copiously produced so the dominant process is A' production through dark bremsstrahlung followed by on-shell mediator decay $A' \rightarrow \chi_1 \chi_2$. The reaction $eN \rightarrow eNA'$, where N is a nucleus of atomic (mass) number $Z(A)$ has been well studied in Refs. [46–48]. The relevant features of the reaction can be better illuminated by considering the calculation using the Weizsacker-Williams (WW) approximation, although we note that all plotted results in this study employ a numerical calculation; see Appendix D for more details of our simulations.

In the WW approximation, the differential production cross section of A' is then given by

$$\frac{d\sigma}{dx d\Omega_{A'}} \approx \frac{4\alpha^3 \epsilon^2 \Phi(q_0, q_f) E^2 x}{\pi U^2} \left[\left(1 - x + \frac{x^2}{2} \right) - \frac{x^2(1-x)m_{A'}^2(E^2\theta_{A'}^2 + m_e^2)}{U^2} \right], \quad (11)$$

where $d\Omega_{A'} = 2\pi d\cos\theta_{A'}$ is the A' solid angle with respect to the beam axis in the lab frame and we have defined

$$U(x, \theta_{A'}) = E^2 x \theta_{A'}^2 + m_{A'}^2 \frac{1-x}{x} + m_e^2 x, \quad (12)$$

$$q_0 = \frac{2U}{E(1-x)}, \quad q_f = m_{A'}. \quad (13)$$

Here E is the electron beam energy, $x = E_{A'}/E$ the fraction of the beam energy carried by the A' , and $\alpha\Phi(q_0, q_f)/\pi$ is the WW photon flux for small-virtuality photons with momentum $-t$ bounded by q_0^2 and q_f^2 [46]. For q_0 values much smaller than the inverse nuclear size $\approx 0.4 \text{ GeV}/A^{1/3}$, we have $\Phi(q_0, q_f) \sim Z^2$ up to an overall logarithmic factor.

Ref. [47] found that for any given x , the angular integral is dominated by angles $\theta_{A'}$ such that $Ex\theta_{A'}^2 \lesssim m_{A'}^2 \frac{1-x}{x} + m_e^2 x$. Using this approximation, we can derive a simpler expression for the differential cross section

$$\frac{d\sigma}{dx} = 4\alpha^3 \epsilon^2 \bar{\Phi}(m_{A'}, E) \frac{x^2 + 3x(1-x)}{3m_{A'}^2(1-x)}, \quad (14)$$

where $\bar{\Phi}(m_{A'}, E) \equiv \Phi(q_0 = m_{A'}^2/(2E), q_f = m_{A'})$. After integrating, the total A' production cross-section scales roughly as

$$\sigma(eN \rightarrow eNA') \approx \frac{4\alpha^3 \epsilon^2}{3m_{A'}^2} \bar{\Phi}(m_{A'}, E) [\log \delta + \mathcal{O}(1)], \quad (15)$$

where $\delta \equiv \min(m_{A'}/E, m_e^2/m_{A'}^2, m_e/E)$.

C. Generic Detection Signals (Electron & Proton Beams)

In our regime of interest, namely $\epsilon^2 \alpha < \alpha_D$, A' decays promptly primarily into the $\chi_1 \chi_2$ system, imparting an order-one fraction of its energy to the daughter particles, which emerge from the target with large boosts. The boosted DM system gives rise to three classes of observable signatures: detector target scattering, χ_2 decays, and missing energy/momentum carried away by the DM system.

1. Detector Target Scattering

The χ_1 produced in the dump travels unimpeded, enters the detector situated downstream of the target/dump, and scatters off of target particles in the detector. It scatters via the reaction $\chi_1 T \rightarrow \chi_2 T$, where T here could in principle be an electron, a nucleon, or even a nucleus. Similarly, any population of surviving χ_2 s produced in the target could downscatter through the reverse reaction $\chi_2 T \rightarrow \chi_1 T$. The cleanest signal occurs when T is an electron, with energy typically above 10s of MeV. The production rate is proportional to ϵ^2 and the scattering rate is proportional to $\epsilon^2 \alpha_D$, so the total yield in this channel will be proportional to $\epsilon^4 \alpha_D$.

Specializing to the case of electron targets in the $E_\chi \gg \Delta$ regime, the approximate differential cross section is [48, 49]

$$\frac{d\sigma}{dE_e} = 4\pi\epsilon^2 \alpha \alpha_D \frac{2m_e E_\chi^2 - f(E_e)(E_e - m_e)}{(E_\chi^2 - m_\chi^2)(m_{A'}^2 + 2m_e E_e - 2m_e^2)^2}, \quad (16)$$

where E_e is the electron recoil energy, E_χ is the incident $\chi_{1,2}$ energy, and we define

$$f(E_e) = 2m_e E_\chi - m_e E_e + m_\chi^2 + 2m_e^2. \quad (17)$$

This cross section is valid up to corrections of order Δ/E_χ and $\Delta/m_{A'}$, both of which are very small compared for the benchmarks we consider throughout (see Figs. 6, 8, and 7). However, our numerical results evaluate the exact expression presented following the derivation in App. B 2.

Neglecting subleading corrections and integrating the recoil energy up to E_χ gives the approximate result

$$\sigma_{\chi e} \approx \frac{8\pi\alpha\alpha_D\epsilon^2 m_e E_\chi}{m_{A'}^4}, \quad (18)$$

so the corresponding scattering probability inside the detector becomes

$$P_{\text{scatter}} = n_e \sigma_{\chi e} \ell, \quad (19)$$

where ℓ is the DM path length through the detector and n_e is the electron number density. There are similar expressions for scattering off detector nucleons and nuclei, but, as we will see, most of the relevant bounds and projections below exploit the electron channel.

2. Decay of excited state

One of the most powerful channels at beam dump experiments is the direct decay of $\chi_2 \rightarrow \chi_1 e^+ e^-$, whose partial width satisfies

$$\Gamma_{\chi_2} \approx \frac{4\epsilon^2 \alpha \alpha_D \Delta^5}{15\pi m_{A'}^4} \quad (20)$$

in the $m_{A'} \gg m_1 \gg m_e$ limit.⁹ The χ_2 is produced in the beam dump and decays in a downstream detector (depicted schematically in Fig. 2), so the signal yield scales as

$$N_{\text{signal}} \propto \epsilon^2 P_{\text{survive}} P_{\text{decay}}, \quad (21)$$

where the product

$$P_{\text{survive}} P_{\text{decay}} = e^{-\Gamma_{\chi_2} d / \beta \gamma} \left(1 - e^{-\Gamma_{\chi_2} \ell / \beta \gamma} \right) \quad (22)$$

is the probability for χ_2 to survive a distance d out to do the detector and decay inside after traversing a path length ℓ with boost factor γ and velocity $\beta = v/c$. As with scattering, the total decay yield scales as $\epsilon^4 \alpha_D$, with ϵ^2 coming from production and $\epsilon^2 \alpha_D$ coming from expanding the exponentials in Eq. (22) assuming the long lifetime limit (see Fig. 4) and using Eq. (20).

To estimate the relative reach of decay and scattering searches at beam dump experiments, it is useful to define

$$R \equiv \frac{P_{\text{decay}}}{P_{\text{scatter}}} \simeq \frac{\Delta^5}{60\pi^2 \gamma E_\chi m_e n_e}, \quad (23)$$

where n_e is the number density of detector electrons. Here we have used the approximate χ -electron scattering cross section from Eq. (18); note that this expression is independent of detector geometry or DM production rate.

For most materials, $n_e \sim 10^{24} \text{ cm}^{-3}$, so we have

$$R \simeq 3 \left(\frac{m_1}{20 \text{ MeV}} \right) \left(\frac{\Delta}{0.1 m_1} \right)^5 \left(\frac{10 \text{ GeV}}{E_\chi} \right)^2, \quad (24)$$

where E_χ represents the typical energy of a χ within the detector acceptance. Note that the $R \sim \Delta^5$ scaling implies that decays dominate the reach for most of the splittings shown in Fig. 6, and that for experiments with lower DM energies (e.g. at LSND where $E_\chi \sim \text{few } 100 \text{ MeV}$) the decay reach dominates the scattering reach by several orders of magnitude due to the lower beam energy and smaller boost factors.

3. Upscatter followed by decay

A third possible signal combines the phenomenology of both scattering and decay. χ_1 can upscatter at the detector,

producing a χ_2 along with a recoiling target T . If the χ_2 produced in the upscatter decays inside the detector via the reaction $\chi_2 \rightarrow \chi_1 e^+ e^-$, and if the electron and positron final states are energetic and separated enough, the final state leads to a signature that is not easily mimicked by environmental or beam-related backgrounds [28, 50, 51]. However, we find that the yields in this channel are subdominant to the other two and we will not consider it further.¹⁰ The consequences of this signal for the case of boosted astrophysical DM, where the different kinematics and the absence of an ϵ^2 production penalty make this channel an attractive background-free option, were investigated in detail in Ref. [22].

4. Missing Energy/Momentum

A unique feature of A' production in electron-beam fixed-target collisions is that the A' is typically radiated with an order-one fraction of the incident beam energy (see Sec. III B and more detailed discussions in [23, 24, 46, 47]). This enables a novel detection strategy where the target is now embedded in the detector, and which is based on comparing the energy and momentum of the beam electron before and after it undergoes dark bremsstrahlung. Fig. 3 illustrates a set-up for this detection strategy.

D. Kinematic Thresholds

There are several kinematic thresholds which influence the detectability of the various signals. Clearly, for $\Delta < 2m_e \approx 1 \text{ MeV}$, the excited state χ_2 can no longer decay via $\chi_2 \rightarrow \chi_1 e^+ e^-$. The only possible decay mode remaining is to χ_1 plus neutrinos or 3γ , both of which are sufficiently suppressed that χ_2 is cosmologically long-lived [52]. In this case the only signals are from the recoiling target in χ_1 upscattering or χ_2 downscattering, along with missing energy/momentum. Otherwise, for detectors with finite energy thresholds and angular resolution, we might require the decay signal $\chi_2 \rightarrow \chi_1 e^+ e^-$ to have an energetic, well-separated $e^+ e^-$ pair satisfying $E_e > E_{\text{min}}$ and $\theta_{e^+ e^-} > \theta_{\text{min}}$. The requirement on Δ to allow a visible χ_2 decay is actually more stringent:

$$\Delta^2 \geq 2m_e^2 (1 + \cos \theta_{\text{min}}) + 2E_{\text{min}}^2 (1 - \cos \theta_{\text{min}}). \quad (25)$$

For example, requiring $E_{\text{min}} = 50 \text{ MeV}$ and $\theta_{\text{min}} = 2^\circ$ as at MiniBooNE, we find that the minimum splitting we can probe is $\Delta \gtrsim 2 \text{ MeV}$.

IV. EXISTING BOUNDS AND PROJECTIONS

In this section we discuss the key features of the various representative experiments we consider and list our kinematic

⁹ For $\Delta > 2m_\mu$, χ_2 can also decay to muons, but for the majority of the parameter space we consider in this paper, only the electron channel is allowed.

¹⁰ If the recoiling target T is not visible, the signal in this channel is identical to the decay signal, but with the added $\epsilon^2 \alpha_D$ penalty of scattering.

cuts used to compute reach curves. The bounds and projections we derive are presented in Figs. 6, 7, and 8.

A. Signals at Proton Beam Dump Experiments

1. LSND

The Liquid Scintillator Neutrino Detector (LSND) [53] was a neutrino experiment at Los Alamos which ran from 1993-1998. The extremely high-luminosity proton beam produced the largest available fixed-target sample of neutral pions, $N_{\pi^0} \sim 10^{22}$. The LSND proton beam had kinetic energy 800 MeV, small enough that η production and dark bremsstrahlung are negligible. Therefore, the DM is produced dominantly from π^0 decays. We modeled the π^0 production at LSND using the GEANT simulation from Ref. [45] and then decayed the pions to DM using Monte Carlo. Due primarily to the large luminosity, we expect DM signal event yields at LSND to dominate in the regions of parameters space where pion decay is kinematically allowed.

The LSND detector was an approximately cylindrical tank of scintillating mineral oil with length $\sim 10\text{m}$ and radius $\sim 3\text{m}$, located $\sim 30\text{m}$ from the beam stop. The close proximity to the beam stop and its off-axis orientation gives a large geometric acceptance of tracks originating at the beam stop. The detector used photomultiplier tubes to detect Cerenkov light emitted by charged particles in the scintillator. For the purposes of our simulations, LSND is only capable of identifying lepton tracks but is blind to nucleons, and electrons are indistinguishable from positrons. We take the electron detection efficiency at LSND to be 19% [25].

We numerically computed the event yields at LSND using DM events that were produced as outlined in Sec. III A, for each of the available signal channels in Sec. III C. Using the techniques outlined in App. D 2, we simulated the χ_1 and χ_2 scattering off electrons, nucleons, and nuclei as well as direct χ_2 decays in the detector. This simulation accounted for the geometric acceptance of the detector as well as kinematic cuts and detection efficiencies for the observable final state particles.

Previous dark matter limits at LSND have been derived using the 55-event background-subtracted limit from the LSND neutrino-electron elastic scattering analysis [25, 54]. To derive the LSND decay constraints for inelastic DM, we interpreted the 55-event limit using χ_2 decay events where the e^+e^- opening angle was too small to distinguish the two leptons, and thus could fake an elastic event passing the cuts described in Ref. [25]. The angular resolution of LSND is 12° for electrons above 20 MeV [53], so we require that the angular separation of the e^+e^- pair is $\theta_{e^+e^-} < 12^\circ$ and that the total energy of the pair lies in the range $18 < E_{e^+} + E_{e^-} < 50$ MeV with either the electron or positron satisfying $\cos \theta < 0.9$. This is conservative, as we have ignored a similar number of events where the electron and positron are well-separated; with access to the full LSND data set, one could search for events where two leptons were seen simultaneously inside the detector and potentially improve the reach.

Indeed, we chose such a conservative prescription because, as we will see, the limits derived from LSND are already dominant for $\Delta > 2m_e$, $2m_1 + \Delta < m_{\pi^0}$.

For the LSND scattering constraints, we again use the 55-event limit applied to any event with *at least* one lepton passing the elastic cuts in the final state. This includes χ_2 down-scattering and χ_1 upscattering with and without χ_2 decay, though the dominant process for all but the largest DM masses is χ_2 downscattering, for which the signal is identical to neutrino-electron elastic scattering. We require a recoiling electron track to be forward-oriented with $\cos \theta > 0.9$ and have energy between $18 < E_e < 50$ MeV. Note that because LSND does not have particle ID, a potential background is neutral-current π^0 production where the photons from the π^0 decay fake electrons.

2. MiniBooNE

MiniBooNE [55] is a proton beam neutrino experiment currently operating at Fermilab with beam energy 8 GeV. It has been previously noted that MiniBooNE is sensitive to light DM [54, 56], and indeed, the first limits from a dedicated DM search have recently been published [57]. At MiniBooNE, the beam energy is large enough that η production and dark bremsstrahlung contribute to DM production, however, contributions from π^0 decays dominate in regions of parameter space where the $\pi^0 \rightarrow \gamma\chi_1\chi_2$ is allowed. Due to the smaller luminosity, the number of pions produced at MiniBooNE is only $N_{\pi^0} \sim 10^{20}$, with the number of etas being further suppressed by a factor of ~ 30 . We used BdNMC [44] to generate a distribution of π^0 and η produced at MiniBooNE and then decayed the mesons to DM using Monte Carlo. To simulate dark bremsstrahlung production at MiniBooNE, we used BdNMC to generate a distribution of on-shell A' and then decayed the dark photons to DM using Monte Carlo. The rate of dark bremsstrahlung production decreases as $m_{A'}$ is increased, but for $m_{A'} \gtrsim 200$ MeV, kinematic thresholds and α_D suppression from off-shell meson decay are significant enough that the dark bremsstrahlung production of DM dominates. A more detailed discussion of the Monte Carlo methods and matrix elements used in these simulations can be found in Appendices C and D.

The MiniBooNE detector also uses scintillating mineral oil, and is approximately spherical with radius $\sim 5\text{m}$, located $\sim 500\text{m}$ from the beam stop. The smaller detector and larger distance from the beam stop means that the geometric acceptance is smaller than LSND, but because the beam energy of MiniBooNE is ten times that of LSND, the DM produced at MiniBooNE has boost factors that are roughly ten times greater, which to some extent compensates for the geometric acceptance as the DM is more collimated. That said, the combination of the smaller proton luminosity and the larger number of background events for a DM search [57] results in the MiniBooNE scattering reach being suppressed by several orders of magnitude compared to LSND. We do not include the MiniBooNE scattering curves in our reach plots as they do

not cover new parameter space.¹¹ MiniBooNE is capable of seeing both nucleon and lepton tracks, but electrons are indistinguishable from positrons and neutrons are indistinguishable from protons.¹² We take the efficiency for lepton and nucleon detection to be 35% [60].

At MiniBooNE, we impose energy and angular cuts similar to those in Ref [57]. We require any lepton track to be forward-oriented with $\cos\theta > 0.99$ and have energy in the range $50 < E_e < 600$ MeV. We also require the electron and positron to have an angular separation of at least 2° [60] and an energy of at least 50 MeV to be visible. We determined the decay reach assuming 95% one-sided Poisson c.l., corresponding to 3 events.

B. Signals at Electron Beam Dump Experiments

1. E137

The E137 experiment at SLAC [26] was designed to search for axion-like particles and with a 20 GeV beam delivering ~ 30 C ($\sim 2 \times 10^{20}$ electrons) of current onto an aluminum target positioned approximately 400 m upstream from an aluminum and plastic scintillator detector. Ref. [49] demonstrated that the existing data from this sample is sensitive to sub-GeV DM if the DM can be produced in the beam dump and scatter elastically off detector electrons.

The null result of the E137 search can be used to place bounds on our inelastically coupled scenario. The DM yield is computed by evaluating the χ flux produced in the beam dump via dark-bremsstrahlung (described in Sec. III B) and considering both DM scattering off electron targets in the detector (described in Sec. III C 1), as well as χ_2 decays inside the detector volume (Sec. III C 2). To comply with the search criteria from Ref. [26], we demand either signal process deposit at least 1 GeV of energy inside the detector's geometric acceptance, and we place a 3-event bound to account for the null result of the E137 experiment.

2. BDX

The proposed BDX experiment at Jefferson Laboratory [28, 48, 50, 51] is a dedicated effort to search for light DM produced and detected in analogy with the procedure at E137.¹³ The setup involves placing a meter-scale CsI scintillator detector behind the beam dump at the upgraded 11 GeV

CEBAF beam. In comparison with E137, BDX has greater luminosity, ($\sim 10^{22}$ electrons on target), a shorter baseline (~ 20 m distance from dump to detector), and a larger detector volume ($50 \times 50 \times 200$ cm³).

In the BDX setup, the boosted $\chi_1\chi_2$ system emerges from the beam dump and can either scatter (see Sec. III C 1) via $\chi_i e \rightarrow \chi_j e$ or the excited state can decay via $\chi_2 \rightarrow \chi_1 e^+ e^-$ (Sec. III C 2) as it passes through the ~ 2 meter long detector. We compute decay and scattering projections using 3 and 10 event yield contours, respectively, for EM energy depositions above 300 MeV.

C. Signals at Electron Momentum/Energy Experiments

1. NA64

The NA64 experiment at the CERN SPS [24], depicted schematically in Fig. 3 (right panel), is sensitive to DM production via dark bremsstrahlung as described in Sec. III B. In this setup, 2×10^9 electrons with 100 GeV energies are fired into an active ECAL target which measures the energy/momentum and triggers on events with large missing energy. In principle, this technique is sensitive to our scenario of interest because the excited state can decay outside the detector, thereby giving rise to missing energy in the ECAL measurement. While we include this discussion for completeness, we find that the existing data sample does not currently constrain new parameter space, so we do not discuss it further.

2. LDMX

The proposed LDMX experiment at SLAC [2, 23] aims to produce DM using the 4 or 8 GeV LCLS-2 electron beam passing through a thin target upstream of a dedicated tracker and ECAL/HCAL system designed to veto SM particles produced in these collisions – the setup is depicted schematically in Fig. 3 (left panel). For our inelastically coupled DM scenario, an A' is produced via dark bremsstrahlung (described in Sec. III B) and decays via $A' \rightarrow \chi_1\chi_2$ followed by a displaced $\chi_2 \rightarrow \chi_1 e^+ e^-$ decay. If this decay occurs behind the ECAL/HCAL system, it can mimic the missing energy signature which LDMX is optimized to observe.

A representative realization of this setup involves $\sim 10^{16}$ electrons with 4 GeV energies passing through a tungsten target of thickness $\sim 0.1X_0$ and emerging from the target with less than 1 GeV of energy. The signal yield scales as

$$N_{\text{signal}} \approx N_e - n_W \sigma(eN \rightarrow eNA') X_0 e^{-\Gamma_{\chi_2} \ell / \beta \gamma}, \quad (26)$$

where $N_e = 10^{16}$ is the number of electrons on target, $X_0 \sim 9$ cm is the tungsten radiation length, n_W is the tungsten number density, $\sigma(eN \rightarrow eNA')$ is the A' production cross section from Eq. (15) and β/γ are the χ_2 velocity and boost factors, respectively. Here the exponential factor is the χ_2 survival probability through a path length of ℓ through the LDMX

¹¹ SBND [58] uses the same beamline as MiniBooNE, but with a smaller detector placed closer to the beam stop. We expect the decay reach to be similar to MiniBooNE for all but the largest mass splittings, but the scattering reach should be enhanced due to the higher detector efficiency [59].

¹² However, MiniBooNE can distinguish photons from leptons [60].

¹³ A similar idea involving lower-energy electron beams installed near large underground neutrino detectors was proposed in Ref. [61].

ECAL/HCAL system. The LDMX missing momentum projections present the 3 event yield contours for various parameter benchmarks.

We note in passing that it may also be possible for both NA64 and LDMX to be sensitive to a promptly decaying $\chi_2 \rightarrow \chi_1 e^+ e^-$ inside their ECAL systems, but the backgrounds for this process are not known and understanding this signal is beyond the scope of the present work.

V. PLOTS & MAIN RESULTS

We now have all the ingredients in place to assess the potential sensitivity of currently proposed fixed-target experiments to discover inelastic DM models. We begin with a brief discussion of existing constraints.

The parameter space of inelastically coupled DM for mass scales beneath $\sim \text{GeV}$ is constrained by precision measurements, B-factories, and previous fixed-target experiments. On the precision front, the anomalous magnetic moment of the muon and electron constrains the interaction strength between the dark photon A' and the SM particles [62]. On the collider front, both LEP and BaBar set a bound for larger values of ϵ . The former arises from the shift in m_Z induced from mixing with A' [63], and the latter from a monophoton and missing mass re-analysis [40, 42, 48].

Some of the strongest constraints from elastic DM arise from E137 [26, 49], an electron beam dump experiment, and LSND [25, 54], a proton beam fixed-target neutrino production experiment. Here, we reinterpret the constraints in terms of coannihilating DM. As discussed in Sec. III C above, there are two qualitatively different signals: a scattering signal, where $\chi_{1,2}$ up- or downscatter at the detector and produce a recoiling target and possibly an $e^+ e^-$ pair, and a decay signal, where χ_2 survives to the detector and decays inside, producing an $e^+ e^-$ pair. The reach of these experiments depends on their ability to distinguish these multiple signals. While E137 is only sensitive to total energy deposits, the angular resolution of LSND means it is potentially sensitive to well-separated e^+/e^- pairs, which can be distinguished from the fake elastic events we used in estimating the sensitivity in this work and could enhance the sensitivity. However, this would require access to the LSND data as this signal of two charged tracks in the detector is not present in any published analysis. These existing constraints are illustrated in Fig. 6 for $\alpha_D = 0.1$, $m_{A'} = 3m_1$, and various values of Δ . For all but the smallest splittings, the combination of LSND and E137 covers a large portion of the thermal target in the 1-100 MeV range. However, for $2m_1 + \Delta > m_{\pi^0}$, DM production through pions is kinematically forbidden, so we see sharp kinematic cutoffs at the pion threshold.

For comparison, Fig. 7 shows the effect of varying our benchmark parameters (each panel varies one detail relative to the top-left panel of Fig. 6) to demonstrate that these benchmarks are conservative and representative of the viable parameter space. In particular, the top row of Fig. 7 shows how the parameter space in the top-left plot of Fig. 6 changes as we increase and decrease α_D while holding all other parameters

fixed. Although there is slightly more viable parameter space for the large value $\alpha_D = 0.5$, this choice is close to the perturbativity limit for abelian dark sectors [65], so we regard our benchmark choice as a representative and conservative value; choosing a smaller coupling excludes *more* parameter space on the y vs. m_1 plane, as we see for $\alpha_D = \alpha$ in the same figure. In the bottom row of Fig. 7, we vary our Δ and $m_{A'}$ benchmarks from Fig. 6. In the bottom-left plot, we show the nearly-elastic case of $\Delta = 0.01m_1$, where the decay signal shuts off and the constraints are dominated by scattering. For comparison, we also show the recent MiniBooNE elastic scattering results [57], for which the beam energy is sufficiently large that the small 1% mass splitting does not affect the reach.

In the bottom-right plot, we show results for a larger hierarchy, $m_{A'}/m_1 = 10$. For a given m_1 , Δ , and α_D , the production rate is decreased as that event now arises from a much heavier A' . If we parameterize the production rates at $m_{A'}/m_1 = 3$ and $m_{A'}/m_1 = 10$ as $N_3 \epsilon^2$ and $N_{10} \epsilon^2$, respectively, the total decay or scattering yield scales as $N_{3,10} \epsilon^4 / m_{A'}^4$. Thus, for a fixed event yield, ϵ scales linearly with $m_{A'}$ but only as $N_{3,10}^{1/4}$. Far from any kinematic boundaries, the sensitivity in $y \propto \epsilon^2 / m_{A'}^4$ improves relative to the thermal target since the scaling with $m_{A'}$ dominates the scaling with $(N_3/N_{10})^{1/4}$. However, the reach at large masses degrades as the A' mass approaches the maximum available energy more rapidly and A' production shuts off.

We now turn to the potential of new proposals to largely cover the entire parameter space motivated by thermal inelastic DM. We focus on three experiments representative of the setups we have previously discussed: MiniBooNE, BDX, and LDMX, which are proton beam dump, electron beam dump, and missing energy experiments, respectively. As discussed in Sec. III, the dominant signal at MiniBooNE is χ_2 decay in the detector whenever it is kinematically allowed. Since MiniBooNE has particle ID [60, 67], electrons can in principle be distinguished from photons, and thus a well-separated e^+/e^- pair and no other activity in the detector is a signal with few irreducible backgrounds. This stands in sharp contrast to the case of elastic DM scattering at MiniBooNE [57], which must always contend with an irreducible neutrino background. Note that the lower boundary of the decay curve is set by the energy threshold and angular resolution according to Eq. (25). We did not attempt a detailed simulation of the ρ resonance, and thus the reach in y of the region very close to $m_{A'} \sim m_\rho$ may differ slightly from what we show. The missing energy signal at LDMX dominates at low masses, while the decay reach for BDX is similar to that of MiniBooNE at high masses. Indeed, while the A' production rate at proton beam experiments is a few orders of magnitude larger than at electron beam experiments, the higher luminosity and beam energy at BDX compensate to give a roughly similar reach. This is advantageous given that different models for the mediator could enhance or suppress production at proton beams compared to electron beams, so the combination of both experiments probes a wide range of models.

Fig. 8 combines the results in this work with the results of Ref. [19] to show the thermal target over a wide range of DM masses. We see that except for a few isolated masses, the

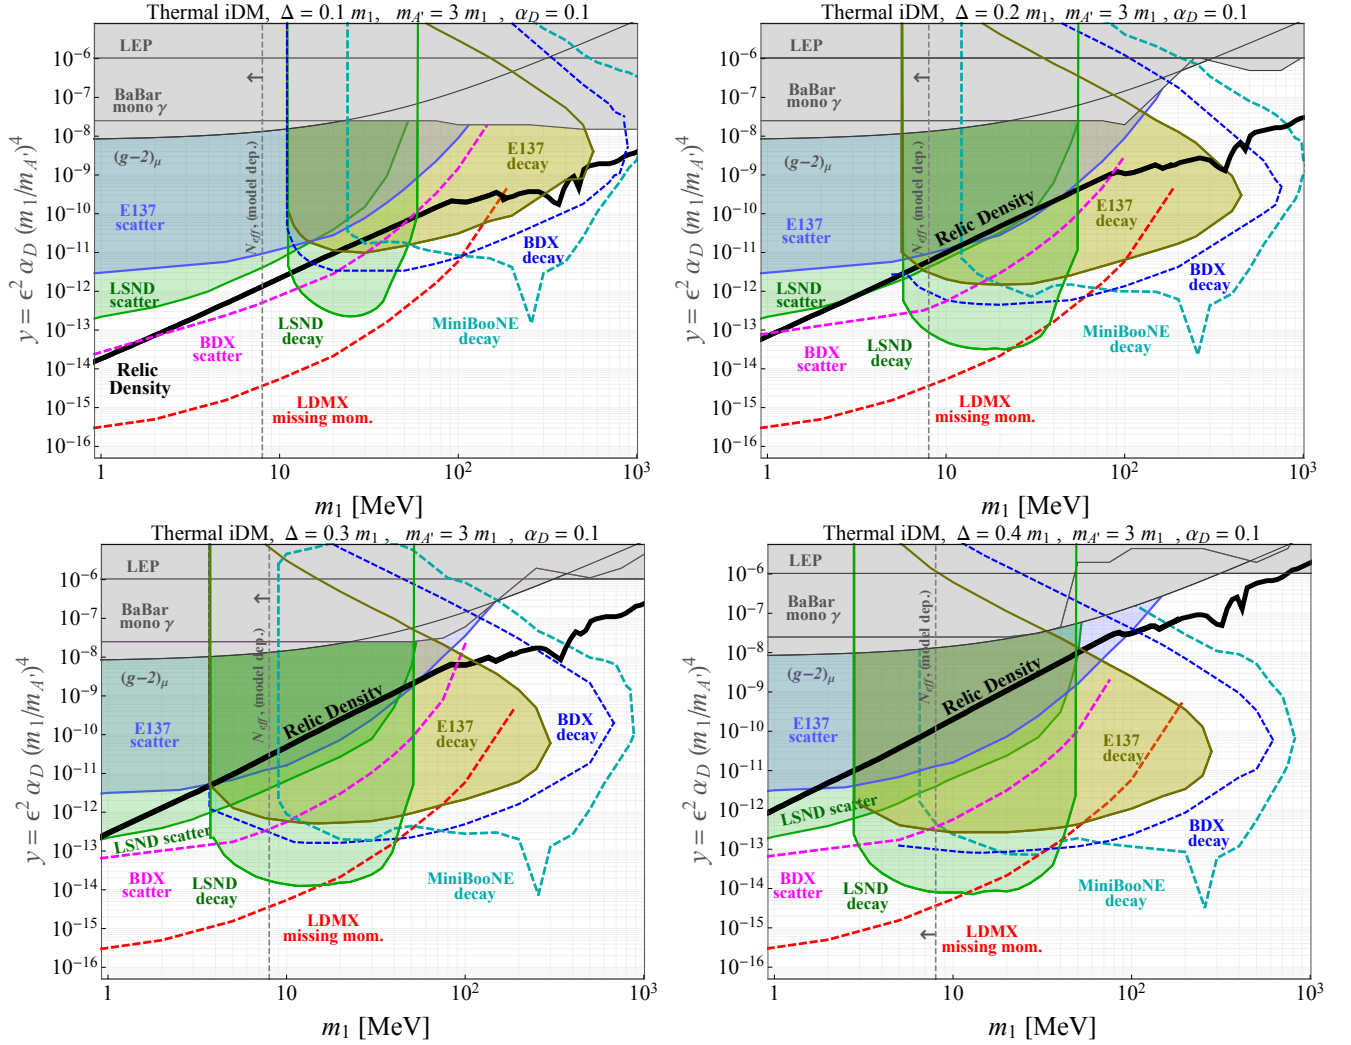


FIG. 6. Parameter space compatible with thermal inelastic DM for different choices of Δ with constraints and future projections presented. The black relic density curve is computed using the procedure described in Appendix A. For each choice of Δ , the relic density curve is insensitive to the relative values of ϵ , α_D , or $m_1/m_{A'}$, however, some constraints depend sensitively on these choices. Typical examples of this sensitivity are LEP and $(g-2)_\mu$ for which the curves shown here are based only on their limits on ϵ ; the observables in question do not depend on α_D or the DM/mediator mass ratio. Thus, where appropriate, we have adopted the conservative prescription $\alpha_D = 0.1$ and $m_{A'}/m_1 = 3$ to place these constraints on this plot, thereby revealing the remaining viable parameter space; see text for a discussion. The colored curves in these plots represent new results computed in work: solid lines are existing constraints, and dashed lines are projections. The gray shaded regions represent kinetic mixing constraints $(g-2)_\mu$ [62]; LEP [63]; and BaBar [19]. Finally, the vertical dashed line labeled N_{eff} is a model-dependent bound from DM freeze-out reheating photons preferentially over neutrinos during BBN [64], excluding parameter space to the left of the line; if there are other sources of dark radiation, this bound can be alleviated.

thermal target for coannihilating DM could be well-covered by all three planned experiments below ~ 1 GeV, and by collider experiments from ~ 1 GeV to 1 TeV. The scattering signals dominate at low mass below the kinematic threshold $\Delta = 2m_e$, while the decay signals dominate when kinematically allowed.

VI. CONCLUSION

In this paper we have studied the fixed-target phenomenology of thermal dark matter with inelastic couplings to the SM

and proposed a series of new searches for these interactions. These models are an instance of the general case where the relic abundance arises from thermal coannihilation between the halo DM candidate χ_1 and an unstable excited state, χ_2 . Since the heavier state decays away in the early universe, there is no annihilation at later times, and therefore no indirect detection. Furthermore, if the mass difference between these states exceeds ~ 100 keV, upscattering at direct detection experiments is kinematically forbidden and loop-induced elastic scattering is small, so this scenario can likely only be discovered or falsified using accelerators. We leave the possibility of one-loop elastic scattering at recently proposed electron direct

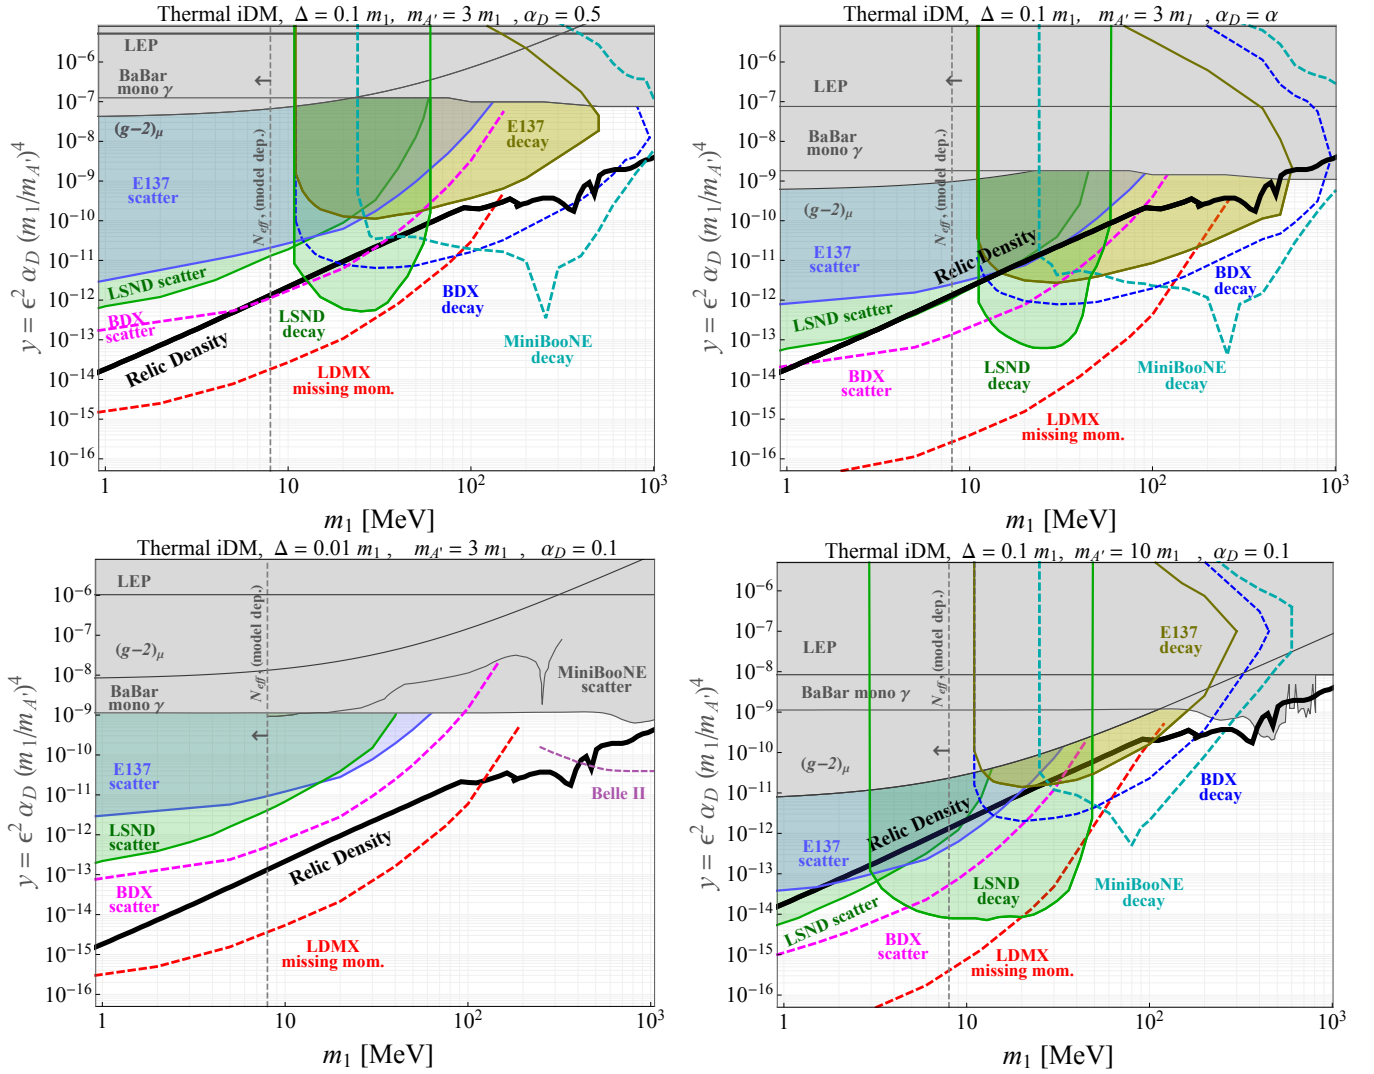


FIG. 7. **Top row:** Same as the top-left panel of Fig. 6, but with different choices for α_D . For larger couplings near the perturbativity limit [65] (left with $\alpha_D = 0.5$) the viable parameter space increases slightly relative to Fig. 6. For smaller couplings (right with $\alpha_D = \alpha$) the thermal target is nearly closed. **Bottom row:** Same as the top left panel in Fig. 6, but with different choices for the inelastic splitting $\Delta = 0.01 m_1$ (left) and the mediator mass $m_{A'} = 10 m_1$ (right). Note that for very small mass splittings, the decay searches become ineffective and the best limits arise from scattering and collider searches, whose observables do not rely on a prompt χ_2 decay.

detection experiments for a future study.

At fixed-target experiments, the inelastic interaction responsible for setting the relic abundance yields a variety of observable signatures arising from the boosted $\chi_1 \chi_2$ system, which is produced in a proton or electron beam collision with target nuclei. Once produced, either state can scatter off particles in a downstream detector, thereby generating an observable signal. In addition, the boosted χ_2 can also survive out to the detector and decay semi-visibly via $\chi_2 \rightarrow \chi_1 e^+ e^-$ to directly deposit a visible signal as it passes through the active volume.

Using these signatures, we have extracted existing constraints on this scenario by reinterpreting old LSND and E137 data. To this end, we have generalized the analyses in [54] (for LSND) and [49] (for E137), which focused on the scattering signatures of elastically coupled DM. In our analysis, we have

demonstrated that there are several new signatures to which these older experiments are sensitive if DM couples inelastically. In particular, we find that E137 and LSND already place nontrivial bounds on the parameter space that yields sub-GeV thermal coannihilation for a variety of DM masses, mass splittings, and coupling strengths.

We have also studied the prospects for future decay and scattering searches at the existing MiniBooNE (proton beam) experiment and the proposed BDX and LDMX (electron beam) experiments. We find that the combined reach of all scattering and decay searches at these experiments can comprehensively test nearly all remaining parameter space, thereby providing strong motivation for these efforts.

This paper also extends earlier work [19], which studied the collider phenomenology of inelastic thermal coannihilation models over the GeV – TeV mass range. Our work com-




FIG. 8. Same parameter space as the top-left and bottom-right panels of Fig. 6, but with the mass range extended out to the electroweak scale. Here we combine the results of this paper with the LHC, BaBar, and Belle II constraints and projections presented in Ref. [19]. The combined reach from the sum of these efforts suffices to cover nearly all remaining parameter space for thermal coannihilation; thermal DM models with masses below the MeV scale suffer generic conflicts with N_{eff} [66] and masses above ~ 100 TeV generically violate perturbative unitarity [5]. The only gaps not covered by this program of searches occur at very small mass splittings $\Delta \ll 0.01 m_1$, depicted in the lower left panel of Fig. 7. For such small splittings, the decay searches become weak on account of the $\Gamma_{\chi_2} \propto \Delta^5$ scaling, and are not even kinematically allowed at low masses since $\Delta < 2m_e$.

plements this effort by working out the constraints and projections for the MeV–GeV range, thereby providing a roadmap for covering thermal coannihilation over nearly all masses for which a thermal origin is viable (lower masses are in conflict with early universe cosmology and higher masses generically violate perturbative unitarity in most models). This full coverage, spanning the results of both papers, is presented in Fig. 8.

Finally we note that other existing and future experiments may also have powerful reach to this class of models. In particular, the proton beam experiments DUNE [68], SeaQuest [69] (see forthcoming work by [70]), SHiP [71], and T2K [72] all involve beam energies in excess of 100 GeV, which can produce far more boosted DM than the beams considered in this work (all < 10 GeV and below). Higher energies at these experiments can open up new production modes for

the DM candidates (e.g. deep inelastic scattering) and impart greater boosts to the DM system, which can profoundly affect the sensitivity projections for these setups. In addition, liquid argon detectors such as ICARUS [73] may be more optimized for seeing the two tracks characteristic of the decay signal.¹⁴ However, working out the implications of these features is beyond the scope of this paper. We look forward to the results of the numerous planned experiments on the horizon, and encourage them to pursue the inelastic DM implications of the signals we have discussed in this work.

¹⁴ We thank Maxim Pospelov for pointing this out.

ACKNOWLEDGMENTS

Acknowledgments: We thank Brian Batell, Asher Berlin, Nikita Blinov, Maxim Pospelov, Philip Schuster, Brian Shuve, Tim Tait, Natalia Toro, Richard Van De Water, and Yim-

ing Zhong for helpful conversations. Fermilab is operated by Fermi Research Alliance, LLC, under Contract No. DE-AC02-07CH11359 with the US Department of Energy. EI is supported by the United States Department of Energy under Grant Contract desc0012704.

-
- [1] G. Jungman, M. Kamionkowski, and K. Griest, Phys. Rept. **267**, 195 (1996), arXiv:hep-ph/9506380 [hep-ph].
 - [2] J. Alexander *et al.* (2016) arXiv:1608.08632 [hep-ph].
 - [3] R. Essig, J. A. Jaros, W. Wester, P. H. Adrian, S. Andreas, *et al.*, (2013), arXiv:1311.0029 [hep-ph].
 - [4] V. Iršič *et al.*, (2017), arXiv:1702.01764 [astro-ph.CO].
 - [5] K. Griest and M. Kamionkowski, Phys. Rev. Lett. **64**, 615 (1990).
 - [6] D. Tucker-Smith and N. Weiner, Phys. Rev. **D64**, 043502 (2001), arXiv:hep-ph/0101138 [hep-ph].
 - [7] K. Griest and D. Seckel, Phys. Rev. **D43**, 3191 (1991).
 - [8] R. T. D’Agnolo, C. Mondino, J. T. Ruderman, and P.-J. Wang, <https://indico.cern.ch/event/394659/contributions/944014/attachments/790207/1083119/Ruderman.2015.pdf>, to appear.
 - [9] P. A. R. Ade *et al.* (Planck), (2015), arXiv:1502.01589 [astro-ph.CO].
 - [10] R. T. D’Agnolo and J. T. Ruderman, Phys. Rev. Lett. **115**, 061301 (2015), arXiv:1505.07107 [hep-ph].
 - [11] R. Essig, A. Manalaysay, J. Mardon, P. Sorensen, and T. Volansky, Phys. Rev. Lett. **109**, 021301 (2012), arXiv:1206.2644 [astro-ph.CO].
 - [12] S. K. Lee, M. Lisanti, S. Mishra-Sharma, and B. R. Safdi, Phys. Rev. **D92**, 083517 (2015), arXiv:1508.07361 [hep-ph].
 - [13] R. Essig, M. Fernandez-Serra, J. Mardon, A. Soto, T. Volansky, and T.-T. Yu, JHEP **05**, 046 (2016), arXiv:1509.01598 [hep-ph].
 - [14] Y. Hochberg, Y. Kahn, M. Lisanti, C. G. Tully, and K. M. Zurek, (2016), arXiv:1606.08849 [hep-ph].
 - [15] T. Emken, C. Kouvaris, and I. M. Shoemaker, (2017), arXiv:1702.07750 [hep-ph].
 - [16] R. Essig, T. Volansky, and T.-T. Yu, (2017), arXiv:1703.00910 [hep-ph].
 - [17] Y. Bai and T. M. P. Tait, Phys. Lett. **B710**, 335 (2012), arXiv:1109.4144 [hep-ph].
 - [18] N. Weiner and I. Yavin, Phys. Rev. **D86**, 075021 (2012), arXiv:1206.2910 [hep-ph].
 - [19] E. Izaguirre, G. Krnjaic, and B. Shuve, Phys. Rev. **D93**, 063523 (2016), arXiv:1508.03050 [hep-ph].
 - [20] P. Schuster, N. Toro, and I. Yavin, Phys. Rev. **D81**, 016002 (2010), arXiv:0910.1602 [hep-ph].
 - [21] D. E. Morrissey and A. P. Spray, JHEP **06**, 083 (2014), arXiv:1402.4817 [hep-ph].
 - [22] D. Kim, J.-C. Park, and S. Shin, (2016), arXiv:1612.06867 [hep-ph].
 - [23] E. Izaguirre, G. Krnjaic, P. Schuster, and N. Toro, Phys. Rev. **D91**, 094026 (2015), arXiv:1411.1404 [hep-ph].
 - [24] D. Banerjee *et al.* (NA64), Phys. Rev. Lett. **118**, 011802 (2017), arXiv:1610.02988 [hep-ex].
 - [25] L. Auerbach *et al.* (LSND Collaboration), Phys. Rev. **D63**, 112001 (2001), arXiv:hep-ex/0101039 [hep-ex].
 - [26] J. Bjorken, S. Ecklund, W. Nelson, A. Abashian, C. Church, *et al.*, Phys. Rev. **D38**, 3375 (1988).
 - [27] I. Stancu *et al.* (MiniBooNE collaboration), (2001).
 - [28] M. Battaglieri *et al.* (BDX), (2014), arXiv:1406.3028 [physics.ins-det].
 - [29] B. W. Lee and S. Weinberg, Phys. Rev. Lett. **39**, 165 (1977).
 - [30] C. Boehm and P. Fayet, Nucl. Phys. **B683**, 219 (2004), arXiv:hep-ph/0305261 [hep-ph].
 - [31] G. Krnjaic, (2015), arXiv:1512.04119 [hep-ph].
 - [32] J. L. Feng, B. Fornal, I. Galon, S. Gardner, J. Smolinsky, T. M. P. Tait, and P. Tanedo, Phys. Rev. Lett. **117**, 071803 (2016), arXiv:1604.07411 [hep-ph].
 - [33] S. Tulin, Phys. Rev. **D89**, 114008 (2014), arXiv:1404.4370 [hep-ph].
 - [34] B. Holdom, Phys. Lett. **B166**, 196 (1986).
 - [35] H. Davoudiasl, H.-S. Lee, and W. J. Marciano, Phys. Rev. **D85**, 115019 (2012), arXiv:1203.2947 [hep-ph].
 - [36] H. Davoudiasl, H.-S. Lee, and W. J. Marciano, Phys. Rev. Lett. **109**, 031802 (2012), arXiv:1205.2709 [hep-ph].
 - [37] H. Davoudiasl, H.-S. Lee, and W. J. Marciano, Phys. Rev. **D89**, 095006 (2014), arXiv:1402.3620 [hep-ph].
 - [38] Y. Kahn, G. Krnjaic, S. Mishra-Sharma, and T. M. P. Tait, (2016), arXiv:1609.09072 [hep-ph].
 - [39] J. D. Clarke, R. Foot, and R. R. Volkas, JHEP **02**, 123 (2014), arXiv:1310.8042 [hep-ph].
 - [40] R. Essig, J. Mardon, M. Papucci, T. Volansky, and Y.-M. Zhong, JHEP **11**, 167 (2013), arXiv:1309.5084 [hep-ph].
 - [41] I. Jaegle (Belle), Phys. Rev. Lett. **114**, 211801 (2015), arXiv:1502.00084 [hep-ex].
 - [42] J. P. Lees *et al.* (The BaBar Collaboration), (2017), arXiv:1702.03327 [hep-ex].
 - [43] S. Teis, W. Cassing, M. Effenberger, A. Hombach, U. Mosel, and G. Wolf, Z. Phys. **A356**, 421 (1997), arXiv:nucl-th/9609009 [nucl-th].
 - [44] P. deNiverville, C.-Y. Chen, M. Pospelov, and A. Ritz, Phys. Rev. **D95**, 035006 (2017), arXiv:1609.01770 [hep-ph].
 - [45] Y. Kahn, G. Krnjaic, J. Thaler, and M. Toups, Phys. Rev. **D91**, 055006 (2015), arXiv:1411.1055 [hep-ph].
 - [46] K. J. Kim and Y.-S. Tsai, Phys. Rev. **D8**, 3109 (1973).
 - [47] J. D. Bjorken, R. Essig, P. Schuster, and N. Toro, Phys. Rev. **D80**, 075018 (2009), arXiv:0906.0580 [hep-ph].
 - [48] E. Izaguirre, G. Krnjaic, P. Schuster, and N. Toro, Phys. Rev. **D88**, 114015 (2013), arXiv:1307.6554 [hep-ph].
 - [49] B. Batell, R. Essig, and Z. Surujon, (2014), arXiv:1406.2698 [hep-ph].
 - [50] E. Izaguirre, G. Krnjaic, P. Schuster, and N. Toro, (2014), arXiv:1403.6826 [hep-ph].
 - [51] M. Battaglieri *et al.* (BDX), (2016), arXiv:1607.01390 [hep-ex].
 - [52] B. Batell, M. Pospelov, and A. Ritz, Phys. Rev. **D79**, 115019 (2009), arXiv:0903.3396 [hep-ph].
 - [53] C. Athanassopoulos *et al.* (LSND Collaboration), Nucl. Instrum. Meth. **A388**, 149 (1997), arXiv:nucl-ex/9605002 [nucl-ex].
 - [54] P. deNiverville, M. Pospelov, and A. Ritz, Phys. Rev. **D84**, 075020 (2011), arXiv:1107.4580 [hep-ph].

- [55] A. A. Aguilar-Arevalo *et al.* (MiniBooNE), Nucl. Instrum. Meth. **A599**, 28 (2009), arXiv:0806.4201 [hep-ex].
- [56] P. deNiverville, D. McKeen, and A. Ritz, Phys. Rev. **D86**, 035022 (2012), arXiv:1205.3499 [hep-ph].
- [57] A. A. Aguilar-Arevalo *et al.* (MiniBooNE), Submitted to: Phys. Rev. Lett. (2017), arXiv:1702.02688 [hep-ex].
- [58] M. Antonello *et al.* (LAr1-ND, ICARUS-WA104, Micro-BooNE), (2015), arXiv:1503.01520 [physics.ins-det].
- [59] R. Van De Water, private communication.
- [60] R. B. Patterson, E. M. Laird, Y. Liu, P. D. Meyers, I. Stancu, and H. A. Tanaka, Nucl. Instrum. Meth. **A608**, 206 (2009), arXiv:0902.2222 [hep-ex].
- [61] E. Izaguirre, G. Krnjaic, and M. Pospelov, Phys. Rev. **D92**, 095014 (2015), arXiv:1507.02681 [hep-ph].
- [62] M. Pospelov, Phys. Rev. **D80**, 095002 (2009), arXiv:0811.1030 [hep-ph].
- [63] A. Hook, E. Izaguirre, and J. G. Wacker, Adv. High Energy Phys. **2011**, 859762 (2011), arXiv:1006.0973 [hep-ph].
- [64] K. M. Nollett and G. Steigman, Phys. Rev. **D89**, 083508 (2014), arXiv:1312.5725 [astro-ph.CO].
- [65] H. Davoudiasl and W. J. Marciano, Phys. Rev. **D92**, 035008 (2015), arXiv:1502.07383 [hep-ph].
- [66] D. Green and S. Rajendran, (2017), arXiv:1701.08750 [hep-ph].
- [67] H.-J. Yang, B. P. Roe, and J. Zhu, Nuclear Instruments and Methods in Physics Research A **555**, 370 (2005), physics/0508045.
- [68] R. Acciarri *et al.* (DUNE), (2016), arXiv:1601.05471 [physics.ins-det].
- [69] K. Nakahara (Fermilab E906/SeaQuest), *Proceedings, 7th International Workshop on Neutrino-nucleus interactions in the few GeV region (NUINT 11): Dehradun, India, March 7-11, 2011*, AIP Conf. Proc. **1405**, 179 (2011).
- [70] A. Berlin, S. Gori, P. Schuster, and N. Toro, to appear.
- [71] S. Alekhin *et al.*, Rept. Prog. Phys. **79**, 124201 (2016), arXiv:1504.04855 [hep-ph].
- [72] Y. Hayato (T2K), *Neutrino physics and astrophysics. Proceedings, 21st International Conference, Neutrino 2004, Paris, France, June 14-19, 2004*, Nucl. Phys. Proc. Suppl. **143**, 269 (2005).
- [73] M. Antonello *et al.*, (2013), arXiv:1312.7252 [physics.ins-det].
- [74] P. Gondolo and G. Gelmini, Nucl. Phys. **B360**, 145 (1991).
- [75] G. Duda, A. Kemper, and P. Gondolo, JCAP **0704**, 012 (2007), arXiv:hep-ph/0608035 [hep-ph].
- [76] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, JHEP **07**, 079 (2014), arXiv:1405.0301 [hep-ph].
- [77] C.-Y. Chen, M. Pospelov, and Y.-M. Zhong, (2017), arXiv:1701.07437 [hep-ph].

Appendix A: Relic Abundance

The relic abundance of χ_1 is governed by a Boltzmann equation whose collision terms involve thermally averaged coannihilation, decay, and inelastic scattering processes. Defining the dimensionless comoving yield as $Y_i \equiv n_i/s$, where $s(T) = 2\pi^2 g_{s,*} T^3/45$ is the entropy density and $g_{s,*}$ is the number of entropic degrees of freedom, the Boltzmann

system can be written as

$$\begin{aligned} \frac{dY_{1,2}}{dx} = & -\frac{\lambda_A}{x^2} \left(Y_1 Y_2 - Y_1^{(0)} Y_2^{(0)} \right) \pm x \lambda_D \left(Y_2 - \frac{Y_2^{(0)}}{Y_1^{(0)}} Y_1 \right) \\ & \pm \frac{\lambda_S}{x^2} Y_f^{(0)} \left(Y_2 - \frac{Y_2^{(0)}}{Y_1^{(0)}} Y_1 \right), \end{aligned} \quad (\text{A1})$$

where $x \equiv m_2/T$ is a dimensionless time variable and $Y_i^{(0)}$ is the comoving equilibrium yield for species i . We define $\lambda_{A,S,D}$, to be the dimensionless annihilation, scattering, and decay rates

$$\lambda_A = \frac{s(m_2)}{H(m_2)} \langle \sigma v (\chi_1 \chi_2 \rightarrow f \bar{f}) \rangle \quad (\text{A2})$$

$$\lambda_S = \frac{s(m_2)}{H(m_2)} \langle \sigma v (\chi_2 f \rightarrow \chi_1 f) \rangle \quad (\text{A3})$$

$$\lambda_D = \frac{\langle \Gamma (\chi_2 \rightarrow \chi_1 f \bar{f}) \rangle}{H(m_2)}, \quad (\text{A4})$$

and we use the Hubble rate $H(T) \simeq 1.66 \sqrt{g_*} T^2 / m_{Pl}$ during radiation domination, where g_* is the number of relativistic degrees of freedom, and $m_{Pl} = 1.2 \times 10^{19}$ GeV is the Planck mass. Solving this system yields the asymptotic value Y_1^∞ at freeze-out near $x \sim 20$, which determines the relic abundance

$$\Omega_{DM} = \frac{\rho_{\chi_1}}{\rho_{cr}} = \frac{m_1 s_0 Y_1^\infty}{\rho_{cr}}, \quad (\text{A5})$$

where $\rho_{cr} = 8.1 h^2 \times 10^{-47}$ GeV⁴ is the critical density and $s_0 \simeq 2969$ cm⁻³ is the present day CMB entropy. An example solution to this system for a representative model point is presented in Figure 9.

Appendix B: Scattering and Annihilation Rates

1. Coannihilation Rate

The amplitude for this $\chi_1(p_1) \chi_2(p_2) \rightarrow f(k_1) f(k_2)$ coannihilation is

$$\mathcal{A} = \frac{i e e g_D}{s - m_{A'}^2} \bar{u}(p_1) \gamma^\mu v(p_2) \bar{u}(k_1) \gamma_\mu v(k_2), \quad (\text{B1})$$

so squaring and averaging initial state spins gives

$$\begin{aligned} \langle |\mathcal{A}|^2 \rangle = & \frac{2e^2 e^2 g_D^2}{(s - m_{A'}^2)^2} \left[(t - m_1 - m_f^2)^2 + (u - m_1 - m_f^2)^2 \right. \\ & \left. + 2m_1 m_2 (s - 2m_f^2) + 2m_f^2 (s - m_1^2 - m_2^2) + 4m_1 m_2 m_f^2 \right]. \end{aligned} \quad (\text{B2})$$

The differential cross section for this process is

$$\frac{d\sigma}{d\Omega_*} = \frac{|\vec{k}_*|}{|\vec{p}_*|} \frac{\langle |\mathcal{A}|^2 \rangle}{64\pi^2 s}, \quad (\text{B3})$$

where $|\vec{p}_*|$ and $|\vec{k}_*|$ are the initial and final momenta in the CM frame and

$$\frac{|\vec{k}_*|}{|\vec{p}_*|} = \sqrt{\frac{(s - m_1^2 - m_2^2)^2 - 4m_1^2 m_2^2}{(s - 2m_f^2)^2 - 4m_f^4}}. \quad (\text{B4})$$




FIG. 9. Example solution to the Boltzmann system in Eq. (A1) for $m_1 = 50$ MeV, $\Delta = 0.3m_1$, $m_{A'} = 3m_1$, and $\langle \sigma v \rangle = 1.7 \times 10^{-23} \text{cm}^3 \text{s}^{-1}$, corresponding to $\Omega_{\chi_1} h^2 = 0.112$ at late times.

Integrating this result, the total cross section becomes

$$\sigma(s) = \frac{|\vec{k}_*|}{|\vec{p}_*|} \frac{8\pi\epsilon^2\alpha\alpha_D s}{(s - m_{A'}^2)^2} \left[(p_1^0 k_1^0)^2 + (p_1^0 k_2^0)^2 + \frac{2}{3} |\vec{p}_*|^2 |\vec{k}_*|^2 + \frac{m_1 m_2}{2} (s - 2m_f^2) + \frac{m_f^2}{2} (s - m_1^2 - m_2^2) + m_1 m_2 m_f^2 \right], \quad (\text{B5})$$

where $p_i^0 = \sqrt{|\vec{p}_*|^2 + m_i^2}$ and $k_i^0 = \sqrt{|\vec{k}_*|^2 + m_f^2}$. Generalizing the derivation in [74] for coannihilation, the thermally averaged cross section is

$$\langle \sigma v \rangle = \frac{1}{N} \int_{s_0}^{\infty} ds \sigma(s) (s - s_0) \sqrt{s} K_1 \left(\frac{\sqrt{s}}{T} \right), \quad (\text{B6})$$

where $s_0 \equiv (m_1 + m_2)^2$, the thermal averaging factor is

$$N = 8m_1^2 m_2^2 T K_2 \left(\frac{m_1}{T} \right) K_2 \left(\frac{m_2}{T} \right), \quad (\text{B7})$$

and $K_{1,2}$ are modified Bessel functions of the first and second kinds.

Note that in our calculation of the relic density, we account for annihilation to hadrons (e.g. $\chi_1 \chi_2 \rightarrow A'^* \pi^+ \pi^-$) by following the procedure described in [19] where the final state phase space is extracted from the measured distribution $R(s) \equiv \sigma(e^+ e^- \rightarrow \text{hadrons}) / \sigma(e^+ e^- \rightarrow \mu^+ \mu^-)$.

2. DM-SM Scattering Cross Section

For pseudo-Dirac DM, with mass splitting $\Delta \equiv m_2 - m_1$, the matrix element for the process $\chi_1(p_1) T(p_2) \rightarrow \chi_2(k_1) T(k_2)$ is given by

$$\mathcal{A} = \frac{\epsilon e g_D}{(t - m_{A'}^2)} [\bar{u}(k_2) \gamma_\mu u(p_2)] [\bar{u}(k_1) \gamma^\mu u(p_1)] \quad (\text{B8})$$

so the squared, spin-average matrix element is

$$\langle |\mathcal{A}|^2 \rangle = \frac{128\pi^2 \epsilon^2 \alpha \alpha_D}{(t - m_{A'}^2)^2} \left[(k_1 \cdot k_2)(p_1 \cdot p_2) + (k_2 \cdot p_1)(p_2 \cdot k_1) - m_1 m_2 (k_2 \cdot p_2) - m_T^2 (p_1 \cdot k_1) + 2m_1 m_2 m_T^2 \right]. \quad (\text{B9})$$

The differential scattering cross section in the CM frame is

$$\frac{d\sigma}{d\Omega^*} = \frac{1}{2\pi} \frac{d\sigma}{d\cos\theta^*} = \frac{\langle |\mathcal{A}|^2 \rangle}{64\pi^2 s} \frac{|\vec{k}_*|}{|\vec{p}_*|}, \quad (\text{B10})$$

where initial/final state momenta satisfy

$$|\vec{p}^*|^2 = \frac{(s - m_T^2 - m_1^2)^2 - 4m_T^2 m_1^2}{4s}, \quad (\text{B11})$$

$$|\vec{k}^*|^2 = \frac{(s - m_T^2 - m_2^2)^2 - 4m_T^2 m_2^2}{4s}. \quad (\text{B12})$$

In terms of lab frame quantities with a stationary target T ,

$$s = (p_1 + p_2)^2 = m_1^2 + m_T^2 + 2m_T E_{p_1}, \quad (\text{B13})$$

$$\begin{aligned} k_1 \cdot p_1 &= -\frac{1}{2} (2m_T^2 - m_1^2 - m_2^2 - 2m_T E_{k_2}) \\ &= E_{p_1}^* E_{k_1}^* - |\vec{p}^*| |\vec{k}^*| \cos\theta^*, \end{aligned} \quad (\text{B14})$$

so we can change variables to obtain the differential recoil distribution

$$d\cos\theta^* = \frac{m_T}{|\vec{p}^*| |\vec{k}^*|} dE_T, \quad (\text{B15})$$

where $E_T \equiv E_{k_2}$ is the targets recoil energy. Thus, we have

$$\frac{d\sigma}{dE_T} = \frac{m_T \langle |\mathcal{A}|^2 \rangle}{32\pi s |\vec{p}^*|^2}, \quad (\text{B16})$$

which serves as an input into all detector scattering calculations for both proton and electron beam dump experiments.

Appendix C: Matrix Elements for DM Production and Detection

1. Meson Decay

The matrix element for pseudoscalar meson decay $m^0(p_1) \rightarrow \gamma(k_1) \chi_1(k_2) \chi_2(k_3)$ ($m^0 = \pi^0, \eta$) is given by

$$\mathcal{A}_{m^0 \rightarrow \gamma \chi_1 \chi_2} = \frac{-i\epsilon e^2 g_D}{4\pi^2 f_{m^0}} \frac{\left(g_{\nu\lambda} - \frac{q_\nu q_\lambda}{m_{A'}^2} \right)}{s - m_{A'}^2 + im_{A'} \Gamma_{A'}} \times \epsilon^{\mu\nu\alpha\beta} k_{1\alpha} q_\beta \epsilon_\mu^*(k_1) [\bar{u}(k_2) \gamma^\lambda v(k_3)], \quad (\text{C1})$$

where f_{m^0} is the meson decay constant, $\Gamma_{A'}$ is the total A' width, $q \equiv p_1 - k_1 = k_2 + k_3$, and $s \equiv q^2$. The spin-averaged square of the matrix element is

$$\begin{aligned} \langle |\mathcal{A}_{m^0 \rightarrow \gamma \chi_1 \chi_2}|^2 \rangle &= \frac{4\epsilon^2 \alpha^2 \alpha_D}{\pi f_{m^0}^2} \frac{1}{(s - m_{A'}^2)^2 + m_{A'}^2 \Gamma_{A'}^2} \\ &\quad \left[(m_{m^0}^2 - s)^2 (s + 2m_1 m_2) - 8s (k_1 \cdot k_2) (k_1 \cdot k_3) \right. \\ &\quad \left. + 2(m_{m^0}^2 - s)(m_2^2 - m_1^2) (k_1 \cdot (k_2 - k_3)) \right]. \end{aligned} \quad (\text{C2})$$

If $m_1 + m_2 \ll m_{A'} \ll m_{m^0}$ then we can reasonably make the narrow width approximation [45] and take

$$\frac{1}{(s - m_{A'}^2)^2 + m_{A'}^2 \Gamma_{A'}^2} \rightarrow \frac{\pi}{m_{A'} \Gamma_{A'}} \delta(s - m_{A'}^2) \quad (\text{C3})$$

The decay width is given by

$$d\Gamma_{m^0 \rightarrow \chi_1 \chi_2} = \frac{1}{4\pi m_{m^0}} \frac{\beta(m_{m^0}^2, 0, s)}{32\pi^2} \frac{\beta(s, m_1^2, m_2^2)}{32\pi^2} \times \langle |\mathcal{A}_{m^0 \rightarrow \chi_1 \chi_2}|^2 \rangle d\Omega_\gamma^* d\Omega_\chi^* ds, \quad (\text{C4})$$

where we have defined the function

$$\beta(s_{12}, s_1, s_2) = \sqrt{1 - \frac{2(s_1 + s_2)}{s_{12}} + \frac{(s_1 - s_2)^2}{s_{12}^2}}. \quad (\text{C5})$$

Here, $d\Omega_\gamma^*$ refers to angles in the m^0 rest frame and $d\Omega_\chi^*$ refers to angles in the $\chi_1 \chi_2$ CM frame.

2. Excited State Decay

The matrix element for $\chi_2(p_1) \rightarrow \chi_1(k_1) f(k_2) \bar{f}(k_3)$ is given by

$$\mathcal{A}_{\chi_2 \rightarrow \chi_1 f \bar{f}} = \frac{\epsilon e g_D}{(s - m_{A'}^2) + i m_{A'} \Gamma_{A'}} \left(g_{\mu\nu} - \frac{q_\mu q_\nu}{m_{A'}^2} \right) \times [\bar{u}(k_1) \gamma^\mu u(p_1)] [\bar{u}(k_2) \gamma^\nu v(k_3)], \quad (\text{C6})$$

where again $\Gamma_{A'}$ is the total A' decay rate, $q \equiv p_1 - k_1$, and $s \equiv q^2$. In this paper we will consider $f = e^-$ exclusively; decay to muons is allowed only for the largest masses and splittings but may provide a distinctive signal at higher-energy experiments. The spin-averaged square is

$$\langle |\mathcal{A}_{\chi_2 \rightarrow \chi_1 f \bar{f}}|^2 \rangle = \frac{16\epsilon^2 e^2 g_D^2}{(s - m_{A'}^2)^2 + m_{A'}^2 \Gamma_{A'}^2} \times [(k_2 \cdot k_1)(k_3 \cdot p_1) + (k_2 \cdot p_1)(k_3 \cdot k_1) + m_f^2(k_1 \cdot p_1) - m_1 m_2(k_2 \cdot k_3) - 2m_1 m_2 m_f^2]. \quad (\text{C7})$$

We note that because we only consider $m_{A'} > m_1$ and $\Delta < m_1$ in this paper, the A' is always off-shell and we never make the narrow width approximation for χ_2 decays.

The decay width is given by

$$d\Gamma_{\chi_2 \rightarrow \chi_1 f \bar{f}} = \frac{1}{4\pi m_2} \frac{\beta(m_2^2, m_1^2, s)}{32\pi^2} \frac{\beta(s, m_f^2, m_f^2)}{32\pi^2} \times \langle |\mathcal{A}_{\chi_2 \rightarrow \chi_1 f \bar{f}}|^2 \rangle d\Omega_1^* d\Omega_f^* ds, \quad (\text{C8})$$

where β is defined as in Eq. (C5), $d\Omega_1^*$ refers to angles in the χ_1 rest frame, and $d\Omega_f^*$ refers to angles in the $f \bar{f}$ CM frame.

3. DM-SM Scattering

The tree-level matrix elements for scattering $\chi_i(p_1) T(p_2) \rightarrow \chi_j(k_1) T(k_2)$ off a pointlike fermionic target have already been computed above in App. B 2. Since

we will also be interested in scattering off targets with substructure such as nucleons and nuclei, we consider the more general scattering process where T is a fermionic target with both a monopole and dipole coupling to electromagnetism. The matrix element for this process is given by

$$\mathcal{A}_{\chi_i T \rightarrow \chi_j T} = \frac{\epsilon e g_D}{t - m_{A'}^2} \left(g_{\mu\nu} - \frac{q_\mu q_\nu}{m_{A'}^2} \right) \times [\bar{u}(k_1) \gamma^\mu u(p_1)] [\bar{u}(k_2) \Gamma^\nu u(p_2)], \quad (\text{C9})$$

where $q \equiv p_1 - k_1$ is the four-momentum carried by the virtual photon and $t \equiv q^2$ is the Mandelstam variable. The Lorentz structure Γ^μ at the target vertex is

$$\Gamma^\mu = F_1(q^2) \gamma^\mu + F_2(q^2) \frac{i q_\nu \sigma^{\mu\nu}}{2m_f}. \quad (\text{C10})$$

Here, $\sigma^{\mu\nu} \equiv \frac{i}{2} [\gamma^\mu, \gamma^\nu]$, and $F_1(q^2)$ and $F_2(q^2)$ are the electric monopole and dipole form factors which depend on the target T . For the purposes of this paper, we take

$$F_1 = \begin{cases} \frac{1}{(1 - q^2/m_p^2)^2} & T = p \\ 0 & T = n \\ Z & T = N \end{cases} \quad (\text{C11})$$

and

$$F_2 = \begin{cases} \frac{\kappa_p}{(1 - q^2/m_p^2)^2} & T = p \\ \frac{\kappa_n}{(1 - q^2/m_n^2)^2} & T = n \\ 0 & T = N \end{cases} \quad (\text{C12})$$

with $\kappa_p \approx 1.79$ and $\kappa_n \approx -1.9$ [54].

The spin-averaged square of the matrix element is

$$\langle |\mathcal{A}_{\chi_i T \rightarrow \chi_j T}|^2 \rangle = \frac{16\pi^2 \epsilon^2 \alpha_D}{(t - m_{A'}^2)^2} \times \left\{ (F_1 + F_2)^2 [t(p_1 + k_1)^2 + 3t(t - \Delta^2) - (m_2^2 - m_1^2)^2] + (F_1^2 + \tau F_2^2) [((p_1 + k_1) \cdot (p_2 + k_2))^2 + (k_2 + p_2)^2(t - \Delta^2)] \right\}, \quad (\text{C13})$$

where $\tau \equiv -\frac{t}{4m_T^2}$.

For an incoming χ_i with the target at rest in the lab frame, the lab-frame differential cross section is then given by

$$\frac{d\sigma_{\chi_i T \rightarrow \chi_j T}}{d\Omega^*} = \frac{1}{2E_\chi} \frac{1}{2m_T} \frac{1}{|\vec{v}_\chi|} \frac{\beta(s, m_f^2, m_f^2)}{32\pi^2} \times \langle |\mathcal{A}_{\chi_i T \rightarrow \chi_j T}|^2 \rangle, \quad (\text{C14})$$

where β is defined as in Eq. (C5).

We note that for the case of coherent scattering off of nuclei ($T = N$), we make the additional insertion of the Helm form factor [50, 75].

Appendix D: Monte Carlo Techniques

Our simulation performs two distinct operations: production of χ_1 and χ_2 pairs and the detection of $\chi_1 \chi_2$ pairs in a detector. Many of the generic techniques used in our simulations such as numerical phase space integration, rejection sampling of differential probability distributions, and computations utilizing detector geometries were borrowed from or influenced by BdNMC [44].

1. DM Production Simulation for Proton and Electron Beams

We use two simulation pipelines, one for proton beams and one for electron beams. Our proton beam production simulation takes as input an unweighted list of four-momenta from one of the DM progenitors that we considered: π^0 , η , and A' from dark bremsstrahlung. The output of the production simulation is an unweighted list of χ_1 and χ_2 particles produced. Each $\chi_1\chi_2$ four-momentum pair in the output list is randomly sampled from the differential decay rate of the progenitor via a rejection sampling method similar to that used in BdNMC.

For electron beam production, we used a modified version of Madgraph@NLO [76] from Ref. [77], which creates a new physics model which contains new particle containers for the nucleus, A' and $\chi_{1,2}$ system. It utilizes elastic and inelastic atomic form factors from Ref. [46]. The elastic component is given by

$$G_{2,el}(t) = \left(\frac{a^2 t}{1 + a^2 t} \right)^2 \left(\frac{1}{1 + t/d} \right)^2 Z^2, \quad (D1)$$

where $a = 111 Z^{-1/3}/m_e$, and $d = 0.164 \text{ GeV}^2 A^{-2/3}$. We also include a quasi-elastic (inelastic) term:

$$G_{2,in}(t) = \left(\frac{a'^2 t}{1 + a'^2 t} \right)^2 \left(\frac{1 + \frac{t}{4m_p^2}(\mu_p^2 - 1)}{(1 + \frac{t}{0.71 \text{ GeV}^2})^4} \right)^2 Z, \quad (D2)$$

where $a' = 773 Z^{-2/3}/m_e$, m_p is the proton mass, and $\mu_p = 2.79$. The general form factor is then

$$\Phi \equiv \int_{t_{min}}^{t_{max}} dt \frac{t - t_{min}}{t^2} (G_{2,el}(t) + G_{2,in}(t)). \quad (D3)$$

2. DM Detection Simulation for Neutrino Detectors

Our detector simulation takes as input the list of $\chi_1\chi_2$ four-momentum pairs from the production simulation. We assume that the various production processes all take place at the beam stop, so that the trajectories of each DM particle is well defined. There are three ways that a $\chi_1\chi_2$ pair produced at the beam stop can be detected: the χ_2 can reach the detector and either decay or downscatter or the χ_1 can reach the detector and upscatter.

For each χ_2 in the input list, a decay length x is randomly selected from the appropriate exponential distribution. If the χ_2 trajectory intersects with the surface of the detector, we

calculate the path length l along the trajectory from the beam stop to the detector. If $x > l$, i.e. if the χ_2 persists until it reaches the detector, then the χ_2 can be detected via direct decay or downscattering. If $x < l$ or if the χ_2 does not intersect the detector, we decay the $\chi_2 \rightarrow \chi_1 e^+ e^-$ via rejection sampling of the differential decay rate. We then process the resulting χ_1 four-momentum in exactly the same way as a χ_1 produced from primary progenitors, but now accounting for the fact that this χ_1 trajectory starts at the point where the χ_2 decayed rather than the beam stop.

Each χ_2 that reaches the detector can either decay or downscatter. We process decay events by weighting by the total probability for the χ_2 to decay inside the detector and performing the decay $\chi_2 \rightarrow \chi_1 e^+ e^-$ via rejection sampling of the differential decay rate. We then process the final state $e^+ e^-$ pair by applying the kinematic cuts described in Sec. IV. We process scattering events by summing over targets and consider scattering off of each target T independently. To avoid double counting with decay events, we make the conservative assumption that the χ_2 can only scatter if it does not decay in the detector. We therefore weight by the probability that the χ_2 persists through the detector and the independent probability that the χ_2 scatters in the detector and perform the scattering $\chi_2 T \rightarrow \chi_1 T$ by sampling from a uniform distribution in the center of mass variables $\cos\theta$ and ϕ , which uniquely specify the final state kinematics. Because we sample our final state kinematic variables from a uniform distribution rather than the true distribution of final states,

$$\mathcal{P}(\cos\theta, \phi) = \frac{1}{\sigma_T} \frac{d\sigma_T}{d\Omega}, \quad (D4)$$

which is proportional to the differential cross section, we must also weight each event by the factor $4\pi\mathcal{P}(x_1, \dots, x_n)$. This weighting scheme enables a cancellation of the total cross section σ_T , making a Monte Carlo computation of σ_T unnecessary and thereby saving significant computational complexity. The cancellation occurs because each event is also weighted by the total probability for scattering which is $\propto \sigma_T$ when Taylor-expanded. Once the scattering final state is sampled and the weights computed, we apply the cuts described in Sec. IV to the recoiling target.

For each χ_1 that reaches the detector, we process the scattering $\chi_1 T \rightarrow \chi_2 T$ using the same method as χ_2 downscattering. Additionally, we weight by the probability that the up-scattered χ_2 will decay in the detector and perform the decay via rejection sampling. The $e^+ e^-$ pair from the de-excitation and the recoiling target are both included when we apply the kinematic cuts described in Sec. IV.